R-4210-AF

ADA 282 557

Estimating Requirements for Aircraft Recoverable Spares and Depot Repair

John B. Abell, Grace M. Carter,
Karen E. Isaacson, Thomas F. Lippiatt

Prepared for the
United States Air Force

RAND

Approved for public release; distribution unlimited

Best
Available
Copy

PREFACE

This report describes the role of management adaptations such as lateral supply, cannibalization, withdrawals of assets from war readiness spares kits, and expedited processing, handling, and transportation in improving logistics system performance in peacetime and wartime. It shows how these management initiatives can reduce the requirement for aircraft recoverable spares, although it does not quantify their costs. It suggests policies and strategies that will reduce the computed spares requirement and increase the effectiveness of the mix of spares procured.

It also discusses the problem of estimating depot component repair requirements and recommends the use of DRIVE (Distribution and Repair in Variable Environments), a computer-based repair prioritization mechanism, as the computing engine for a repair requirements estimation system. Extensions to that system are suggested to take explicit account of repair capacity constraints, budgetary requirements, the projected operating position of the stock fund, and estimated capability in terms of peacetime and wartime aircraft availability goals.

The principal thrust of this research was to enhance understanding of the implications for requirements estimation of demand uncertainty and logistics management adaptations to cope with it. Several other reports describe the larger body of work of which this is a part and are listed here:

- John L. Adams, John B. Abell, and Karen E. Isaacson, *Modeling and Forecasting the Demand for Aircraft Recoverable Spare Parts,* RAND, R-4211-AF/OSD, 1993.
- John B. Abell and Frederick W. Finnegan, *Data and Data Processing Issues in the Estimation of Requirements for Aircraft Recoverable Spares and Depot Repair,* RAND, MR-264-AF (forthcoming).
- Donald P. Gaver, Karen E. Isaacson, and John B. Abell, *Estimating Aircraft Recoverable Spares Requirements with Cannibalization of Designated Items,* RAND, R-4213-AF, 1993.
- Karen E. Isaacson and Patricia M. Boren, *Dyna-METRIC Version 6: An Advanced Capability Assessment Model,* RAND, R-4214-AF, 1993.

- John B. Abell, *Estimating Requirements for Aircraft Recoverable Spares and Depot Repair: Executive Summary*, RAND, R-4215-AF, 1993.

The first of these reports describes improved methods for forecasting the demand for aircraft recoverable spares and specifying the variance of the probability distribution describing the number of assets of a given type in resupply (i.e., in base repair, retrograde, depot repair, serviceable shipment, or the condemnation pipeline). The second report discusses data and data-processing issues related to the estimation of aircraft recoverable spares and repair requirements. The third presents a computational algorithm for estimating requirements for aircraft recoverable spares based on the assumption that items can be designated as cannibalizable or not. The fourth describes DynaMETRIC Version 6, the capability assessment model used to evaluate stockage postures that were anticipated to eventuate from purchases of particular mixes of recoverable spares. The fifth summarizes the entire body of work.

This work had the joint sponsorship of Headquarters, United States Air Force (AF/LEX), and Headquarters, Air Force Materiel Command (AFMC/XP and AFMC/XR). It was carried out in the Resource Management and System Acquisition Program of Project AIR FORCE, RAND's federally funded research and development center supported by the U.S. Air Force. It should be of particular interest to those concerned with spares and repair requirements estimation, logistics system design and modeling, and logistics policy analysis. It should also interest logisticians throughout the Air Force, the other military services, and in the Office of the Secretary of Defense.

Accession For
NTIS GRA&I ☑
DTIC TAB ☐
Unannounced ☐
Justification
By
Distribution/
Availability Codes
Dist | Avail and/or Special
A-1

SUMMARY

Estimating requirements for aircraft recoverable spares and depot-level repair is a difficult problem, involving substantial budgetary resources. In the mid to late-1980s the Air Force spent roughly $5 billion annually for these resources, about $3 billion for spares procurement and $2 billion for repair. It now faces dramatic budgetary reductions that present a formidable challenge to resource managers. The spares and repair requirements estimation system is large and complicated, involving five air logistics centers and tens of thousands of stock numbers, and is fraught with uncertainty.

The research described in this report involved extensive analysis and evaluation of the current system. Its principal goal was to understand better the implications of management adaptations for spares requirements. By management adaptations we mean such initiatives as cannibalization; lateral supply; withdrawals of assets from war readiness spares kits (WRSK); and expedited repair, processing, handling, and transportation. These and other management initiatives enhance the performance of the logistics system in the face of uncertainty in resource demands and as item characteristics evolve over time, but they are not now accounted for in the computation of spares requirements.

Our overall conclusion is that the Air Force could achieve satisfactory levels of aircraft availability with substantially less expenditure on spares procurement by taking explicit account of the payoffs of management adaptations, by making certain improvements to the computations and processes involved in determining spares requirements, and by implementing certain policy changes. We describe those improvements and policy changes in this report and, to the extent possible, estimate their effects on requirements estimation and system performance. On the other hand, we do not estimate the costs of implementing those changes.

THE ROLE OF UNCERTAINTY IN LONG SUPPLY

The Air Force has been criticized in recent years for having too many items in "long supply," i.e., in an excessively rich asset position. We define long supply simply as assets in excess of the number required a lead time beyond the buy point as estimated by the spares requirements computation. The criticism about long supply is only partially

justified simply because the problem of long supply has its roots in the substantial uncertainty associated with parts demand processes and changing item pipelines. The probability that an item will be in long supply, as defined here, at any point during its life in the inventory system is an increasing function of its pipeline and the variability in its underlying demand process. Parts demand processes are frequently nonstationary, i.e., the expected number of demands for a given stock number in a time period of specified length varies over time, and the magnitude and direction of that variation are almost never predictable. The problem of minimizing the number of items in long supply given a system performance goal is made extremely difficult by this uncertainty. Obviously, retention policy also plays an important role. The role of uncertainty in exacerbating the long supply problem is illustrated in this report.

Some amount of long supply is a natural consequence of uncertainty. It cannot be eliminated, only attenuated somewhat by the use of hedging strategies and lower investment levels. Several of the suggestions we make in this report can help mitigate the effects of uncertainty on the long supply problem.

DEMAND MODELING AND FORECASTING

The levels of uncertainty that pervade the system, coupled with the long procurement lead times involved in spares purchases, make the problem of forecasting demands very difficult. Our work on the forecasting problem suggests that improvements can be made, but those improvements do not solve the problem completely. We do recommend two improvements to demand modeling and forecasting that are discussed in detail in Adams et al. (1993). The first is the use of weighted regression forecasting of demands as the default method in requirements estimation. The second is an improved specification of the variance of the probability distribution of the number of items of each type in resupply, i.e., in base repair, retrograde, serviceable transit, depot repair, or the condemnation pipeline. These improved approaches reduce predictive error roughly 40 to 50 percent. Table S.1 reflects the results of our evaluations of these methods with Dyna-METRIC Version 6 using a replication of the fiscal year 1987 spares requirements computation (using the March 1986 database). The use of the improved methods of forecasting and variance specification resulted in an estimated cost reduction of almost a quarter of a billion dollars ($239 million) in the requirements computation done with the March 1986 database while achieving a level of performance

Table S.1

Performance of Improved Techniques Compared to Current System

Management Adaptations	Percentage of Aircraft Unavailable, Peacetime	
	Current System, $3709 Million	Improved Methods, $3470 Million
No cannibalization	74.9	76.3
Full cannibalization	33.0	33.1
Cannibalization, lateral supply	17.3	17.2
Cannibalization, quick, lateral supply	3.2	3.6

almost equal to the current system. War readiness spares were not included in these evaluations.

UNCERTAINTY AND MANAGEMENT ADAPTATIONS

In the following discussion, we evaluate the effects of several management adaptations including one we call *quick response*. What we mean by quick response is the implementation of initiatives to reduce the processing, handling, and transportation times in the logistics system. The magnitude of reduction we assume is defined in Table S.2. Note that quick response reduces the total depot repair turnaround time by a little less than half. It has powerful effects on system performance. Although evaluation of the costs of achieving such reductions is beyond the scope of this work, it seems fairly clear to us that the payoffs from these adaptations would exceed their costs significantly.

Table S.2.

The Quick Response Option

	Pipeline Lengths	
	Actual	Quick
Base processing days	4[a]	3
Retrograde days	16	5 overseas, 2 CONUS
Supply to maintenance	3[a]	1
Shop flow days	40[a]	40[a]
Serviceable turn-in days	5[a]	2
Order-and-ship time	21	5 overseas, 2 CONUS
Total	89[a]	56[a] overseas, 50[a] CONUS

[a]Varies by item; average shown.

Our evaluations of the current system were done by computing the spares requirement just as the current system does. We replicated the software with which AFMC carries out the Central Secondary Item Stratification (CSIS), and acquired the Aircraft Availability Model (AAM) through the courtesy of the Logistics Management Institute, its developer. The AAM computes requirements for safety stock, which the CSIS treats as an additive. The current asset position, together with the spares buy, enables us to estimate the total number of assets of each type that will be in the inventory system an average lead time beyond the buy point, roughly three years in the future. We built a replica of AFMC's Central Stock Leveling System, D028, and used it to allocate those assets to the bases and depot. The resulting allocations constitute what we refer to as the *anticipated stockage posture.* Our approach to the evaluations was to compute spares requirements worldwide and to evaluate the performance of the anticipated stockage posture with Dyna-METRIC Versions 4 and 6, using the F-16 weapon system (all series) as a case study. The models estimate aircraft availability as a function of asset position. Version 4 is incorporated in the Weapon System Management Information System (WSMIS) maintained by Headquarters, AFMC. Version 6 is an enhanced version.

Figure S.1 shows the results of our evaluations of the performance of the stockage posture anticipated to result in about mid-1989 as a result of the spares requirements computation done with the March 1986 D041 database.[1]

The leftmost dark grey bar in Figure S.1 reflects the aircraft availability goals specified to the requirements computation for the F-16 aircraft: 83 percent availability, or not more than 17 percent unavailable. The second dark grey bar quantifies estimated system performance using the same assumptions about the logistics system that are built into the requirements estimation system: no cannibalization, unspecified lateral supply, and future demand rates and pipeline times the same as in the past. It is higher than the goal because of a limitation in our available data. Without aircraft configuration data, we are forced to assume that items whose application percentages are less than 100 percent (i.e., they are not installed on every aircraft of a specified model-design-series) have the same application percentage at every base with F-16 aircraft because we do not have base-specific

[1]These evaluations used variance-to-mean ratios in all cases that correspond to the formula used in the requirements computation. This is explained in Section 2.

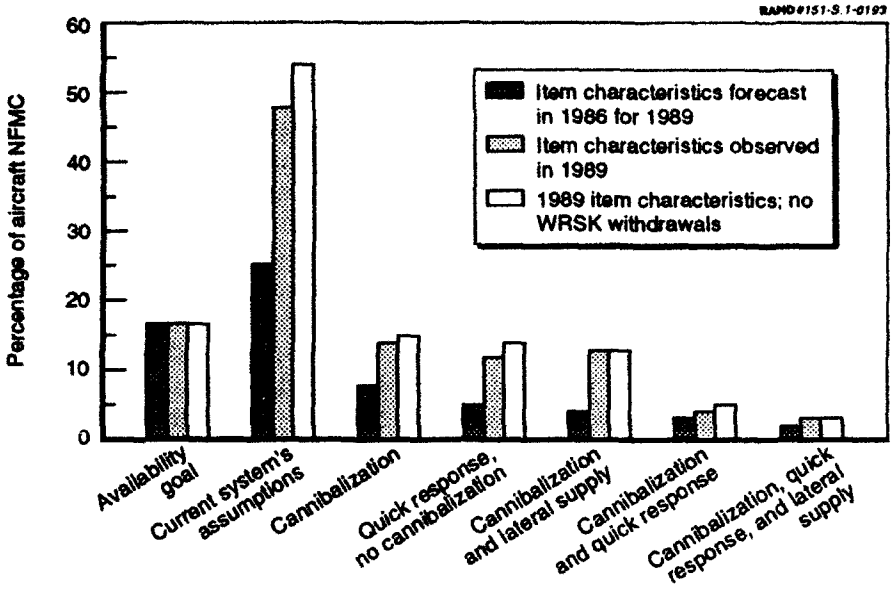

Figure S.1—Effects of Uncertainty, Management Adaptations, and Prohibited WRSK Withdrawals on Peacetime Performance

application percentages. (The computation of base-specific application percentages would require a mapping of stock number into aircraft tail number, coupled with a force beddown by tail number and base.) Our estimates of performance, therefore, do not account for the economies of scale that would result from knowing the actual distribution of expected demands among bases because items with less than 100 percent application are often used at only a fraction of the bases. The fewer bases, therefore, often have higher expected demands than we assume in our evaluation models.

The remaining dark grey bars in Figure S.1 tell one story of compelling interest here: cannibalization, lateral supply, and the expedited actions described as elements of the quick response option dramatically improve aircraft availability. In fact, with all three of these management adaptations in place, aircraft availability reaches 98 percent, substantially exceeding the goal specified to the requirements computation, even in the face of the naive distribution assumption we are forced to make.

The evaluations represented by the dark grey bars in Figure S.1 were done with the same assumptions made in the spares requirements computation. Thus they assume perfect information in the sense that

the demand rates, not-repairable-this-station (NRTS) rates, pipeline times, and other item characteristics used to estimate the performance of the stockage posture anticipated to evolve in 1989 were exactly those reflected in the March 1986 requirements database used to compute the requirement. The light gray bars in Figure S.1 tell the second compelling story in these data. They show the results of evaluating the anticipated stockage posture using item demand rates that *actually eventuated in 1989*. Thus they reflect the effects of uncertainty and *churn* in the database. We define churn as the sum of all changes in the database from one point in time to some future time; item characteristics of the variety already mentioned change over time as Air Force experience evolves and changes, new items appear in the database, and some items that are in a particular database are no longer there at some future time. Repair and transportation times may also change, and the performance of the anticipated stockage posture is vulnerable to this churn. Its total effect is portrayed in Figure S.1 by the light gray bars, which are uniformly higher than the dark gray bars.[2] The third compelling story, also illustrated by the light gray bars, is that despite the formidable effects of churn on system performance under the naive assumptions made in the requirements system shown by the bars labeled "Current system's assumptions," these several management adaptations almost completely overcome those effects, and with all three of them in place, as shown in the rightmost set of bars, aircraft availability still far exceeds the specified goal of not more than 17 percent unavailable.

In summary, then, churn and other uncertainties degrade system performance, but management adaptations have significant offsetting effects. On balance, management adaptations yield levels of system performance well beyond the level specified to the spares requirements computation. This is not to suggest that the wrong goal is being used. The Air Force's "real" goal may be 95 percent aircraft availability. In fact, during this period of time we observed not-mission-capable, supply (NMCS) rates of about 9 percent for the F-16 force. The point is that, when we repeatedly specify an 83 percent goal to the requirements computation in the face of observed aircraft availability rates in the 90s, we are implicitly specifying the higher rates for which the number 83 is only a numerical surrogate.

[2]The variance of the number of assets in resupply was specified to follow the formula used by AFMC in computing the requirement. Performance would have been considerably worse had we used the observed demand variance.

CONSOLIDATING WAR READINESS SPARES

The current policy of authorizing every squadron with wartime deployment tasking a war readiness spares kit (WRSK) derived from consideration of wartime scenarios of the past that are no longer of significant concern. In the current world, deployment contingencies are imagined to be of quite different scope than was the NATO scenario. It makes sense, given the magnitude and implications of recent changes in global threats, to rethink the policy governing the allocation of war readiness spares. An alternative to the current WRSK policy is examined in this work, that of consolidating WRSK at one or a very few locations.

The third set of bars in Figure S.1, the white bars, reflect the results of denying the bases access to war readiness spares in peacetime. The modest effects of this policy may be somewhat surprising because WRSK withdrawals are frequently used to alleviate parts shortages affecting mission capability (MICAPs). This is so because a WRSK withdrawal, although involving some paperwork and permissions, is often more convenient than cannibalization with its higher cost and attendant risk of damage. In any event, these evaluations show a modest effect of holding war readiness spares kits inviolate during peacetime, a policy that might be implicit in consolidating war readiness spares kits in one or a few central locations where they could be computed, assembled, and shipped promptly in response to deployment contingencies. On the other hand, a consolidated WRSK policy could be implemented without the constraint of denying bases any access to peacetime use of war readiness spares; in this case, of course, lateral supply would also be involved. An aircraft configuration database would also be needed to support this concept, but we estimate that consolidation of WRSKs would have substantial payoffs in reduced manpower at the bases and substantially reduced spares requirements. Even if we bought enough war readiness spares to support an operation twice as large as Operation Desert Shield/Storm, we could save as much as $327 million for the F-16 weapon system alone.

ENHANCING THE RESPONSIVENESS OF DEPOT REPAIR

One implication of these evaluations is especially persuasive: Reducing depot component repair turnaround time has dramatic effects on system performance. Given the Air Force's plans to extend the application of the concept of two levels of maintenance, shortening the depot repair pipeline times of selected items seems imperative. An aggressive, continuing program of initiatives should be pursued to ensure that the most relevant mix of assets is being repaired in a

timely manner, and initiatives are being pursued to enhance the robustness of the depot's performance in the face of uncertain demands for repair parts and repairable carcass generations.

DESIGNATED CANNIBALIZATION

The concept of designating items as "cannibalizable" or not in the WRSK database could be extended to primary operating stock (POS). The idea of designating cannibalizable items in the database is attractive because it enhances the effectiveness of the spares mix and helps minimize the investment required, an especially attractive strategy in times of serious budgetary constraints. The ability to designate items that are easy to cannibalize considerably weakens the argument underlying the traditional Air Force policy that cannibalization be ignored in the computation of POS requirements.

We were unable to evaluate this initiative as objectively as we would have liked, but we did try to approximate its effects by somewhat arbitrarily designating all the line-replaceable units (LRUs) and shop-replaceable units (SRUs) in certain stock classes as cannibalizable. The designated items constituted about 58 percent of the total cost of all item pipelines. Figure S.2 shows the results. The same aircraft availability goal was specified to the requirements computation in each case.

System performance remains about the same with designated cannibalization as it is with the current system's assumption of no cannibalization, but the requirement is reduced by $247 million. These results, while not tied to actual judgments of maintenance experts, suggest that the payoffs of the policy of designating cannibalizable items in the requirements database are very attractive.[3]

ESTIMATING REPAIR REQUIREMENTS

Section 6 of this report presents the results of a detailed analysis of component repair requirements for certain F-16 avionics components computed under the current procedure and with DRIVE (Distribution and Repair in Variable Environments) (Abell et al., 1992; Miller and Abell, 1992). The results of this analysis contrast the two approaches and examine detailed data describing the worldwide asset positions of

[3]These evaluations used variance-to-mean ratios actually observed in 1989, but capped at 15.0 to exclude data errors. About 5 percent of the items were affected by the cap.

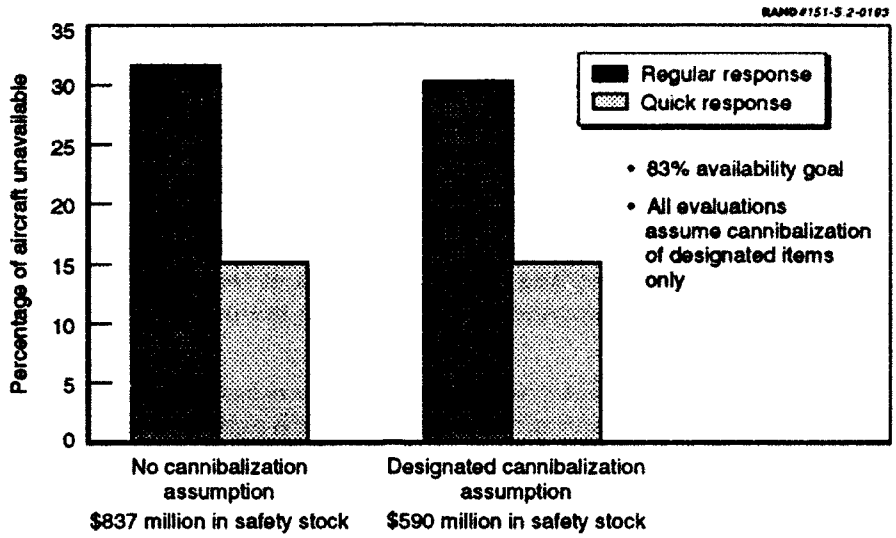

Figure S.2—Designated Cannibalization Pays Off

items for which DRIVE estimated substantially different repair requirements than those reflected by the X21 scrubbed requirements. Our conclusion is unambiguous: DRIVE consistently estimated repair requirements that are much more sensibly tied to the asset position at the bases and to expected NRTS actions than did the current system. It also produced a critical item list that was highly correlated with the Tactical Air Command (TAC, now the Air Combat Command) critical item list published at about the same time. We conclude that DRIVE should be the computing engine of whatever depot component repair planning and management system evolves in AFMC. However, if it is to become a viable system, intensive effort must be devoted to cleaning up and maintaining the data that DRIVE uses in determining repair requirements.

TRADEOFFS AMONG SPARES, REPAIR, AND TRANSPORTATION

System-level tradeoffs among spares, repair, and transportation were beyond the scope of this work. We note, however, the clear direction of the tradeoff: Give up some spares investments in favor of funding initiatives to reduce depot repair turnaround times. Although this may be a difficult strategy to implement because of the partitioning of management responsibilities in the logistics system, it has large payoffs.

OTHER INITIATIVES

Several other improvements to the spares and repair estimation process are developed and discussed at length in this report. Among them are the development of an aircraft configuration database and means to update it routinely, the addition of engines and engine modules to the requirements database to obviate the need to promote engine components to LRU status, improvements in data specification and maintenance, and adding SRU safety stock at the depot. Each of these suggestions is described in detail in the main body of this report.

The discussion that follows summarizes our conclusions and recommendations. They are also discussed at greater length in Section 8.

CONCLUSIONS AND RECOMMENDATIONS

We are impressed with the difficulty of the problem of estimating spares requirements. In general, we conclude that the system overinvests in spares and that the mix of spares bought could be more cost-effective. This work has left many technical issues unresolved, in part because we lack data, and in part because our evaluative tools sometimes fail us. Additional research is needed to understand these technical issues better and to evaluate costs. Some of our recommendations suggest immediate implementation of certain initiatives; some suggest further research and evaluation.

We recommend that AFMC:

- Implement a massive, intensive effort to correct the D041 application file, and a continuing training program in principles and procedures for specifying elements of data that describe item characteristics and applications and in file maintenance. This training program should be coupled with a system of audits of LRU families that are sampled from the population in a way that accounts for item managers, equipment specialists, and commodity groups. The audits should be done to find data problems and identify item managers and equipment specialists who need assistance or training. Results should be analyzed to identify systemic problems.
- Establish an aggressive and continuing program to enhance the responsiveness of depot component repair and the management of all segments of the depot repair pipeline, and couple those functions more closely to the combat force. Cost analysis should be a first step.

- Implement the improved demand forecasting and variance specification techniques described in Adams et al. (1993) as the default methods for high-demand items, and reduce investment levels to maintain desired system performance.
- Implement an enhancement to the Aircraft Availability Model (AAM) such as that described in Gaver et al. (1993) to support the policy of designated cannibalization, and modify the requirements database to reflect the designation of selected items as cannibalizable.
- Add engines and engine modules to the requirements database.
- Develop a configuration database for all aircraft that contains a mapping of subgroup master stock numbers to aircraft serial numbers, and a system for routinely updating it to reflect changes to the mapping.
- Provide for SRU safety stocks for the depot repair system in the spares requirements computation. Again, costs should be estimated as a first step.

The suggestions that follow depend on the aircraft configuration database being in place:

- Consolidate the storage and management of war readiness spares, either at a single location, at a location for each command, or at a location for each weapon system, and develop the ability to compute, assemble, and deploy war readiness spares kits in a very short time (say 24 hours), on very short notice (say 6 hours). This capability should include mission support kits, enroute support kits, and similar contingency requirements.
- Compute quarterly component repair requirements using DRIVE as the computing engine. Extend such requirements computations over longer planning horizons as needed for resource allocation and capital investment decisionmaking.
- Evaluate the cost-effectiveness of base-specific data coupled with a model of lateral supply in the requirements computation. By "base-specific data" we mean (1) the actual number of bases at which items are exposed to OIM demand rather than the notional *number of users* used in the current system, and (2) base-specific application percentages, expected demands, aircraft availability goals, and stock-level allocations.

The effects of modeling lateral supply and using base-specific data tend to be offsetting. Combining them, we believe, would yield a more

effective mix of spares. The use of base-specific data would make the calculations of the AAM considerably more tedious unless the estimation of depot stock levels were partitioned from that of base stock levels. Modeling lateral supply could also help alleviate the problem, depending on how it was done.

Our remaining suggestions are more technical. Although they may not yield large savings, they are important, and may be very important, to the effectiveness of the mix of spares that emerges from the requirements computation:

- Undertake further analysis to resolve the major differences between D041 item programs and those implied by the K004 and application files.
- Evaluate the cost-effectiveness of adding bench-check-serviceable (BCS) demands to item pipelines.
- Delay the specification of the procurement quantities of spares being purchased as long as possible before obligating the government to the contract. Use this additional time to observe additional demand and NRTS data to revise the procurement quantity. (We believe that this initiative is already in place.) Couple this strategy with specification of procurement quantities that leave some or all of additional safety stock requirements as an option for future procurement.

We believe that these several suggestions would result in procurement of more effective mixes of aircraft recoverable spares as well as substantial reductions in spares investments. Implementing improved demand forecasting and variance specification, reducing the investment level to maintain desired system performance, consolidating WRSK at, say, twice the size needed to support Operation Desert Shield/Storm, and implementing the designated cannibalization concept would have reduced the stock levels that emerged from the March 1986 computation by roughly a billion dollars. We do not understand the implications of implementing all of the suggestions described in this section together. Thus we have tried to distinguish ideas that we think should be implemented without further evaluation from those whose effects we can only estimate given the limitations of our models and data.

We realize, too, that we are being somewhat imprecise about an implementation strategy, but our priorities are implicit in the list of recommendations, except for the last three, which can be undertaken immediately.

In summary, however, we are persuaded that the Air Force can achieve roughly its traditional levels of system performance at substantially reduced spares investment costs.

ACKNOWLEDGMENTS

We are indebted to our RAND colleague, Frederick W. Finnegan, for his competent and extensive data-processing support. His intimate knowledge of D041 data files and other data systems was invaluable. James S. Hodges and Stanley S. Bentow of RAND carried out the analysis of long supply discussed in Section 3. We are also grateful to Victor J. Presutti, Jr., Director of Management Sciences at Headquarters, AFMC, for his suggestions, comments, and support, and for sharing with us his extensive knowledge of the requirements estimation process. Colonel Douglas J. Blazer, USAF, who was assigned to the Air Staff at the time this work began, also provided helpful comments and suggestions. We are grateful to the Logistics Management Institute for providing us with a copy of the Aircraft Availability Model and many of the data files on which the analyses presented here were based.

We owe a special debt to Louis W. Miller of RAND for his thoughtful and constructive review of this manuscript. He also contributed the idea of using the geometric mean as a performance measure of worldwide asset positions of LRU families. Much of our discussion of the estimation of component repair requirements is based on his imaginative involvement with this issue over the past several years. Mary Chenoweth of RAND also provided helpful comments.

CONTENTS

PREFACE iii

SUMMARY v

ACKNOWLEDGMENTS xix

FIGURES xxiii

TABLES xxv

GLOSSARY xxvii

Section

1. INTRODUCTION 1
   - Stock Funding of Recoverables 3
   - Force Drawdown 4
   - What Follows 4

2. A BRIEF OVERVIEW OF SPARES REQUIREMENTS ESTIMATION 6
   - Some Common Terms 6
   - The Concept of Item Pipelines and Resupply 7
   - The Role of Uncertainty in Spares Requirements Estimation 9
   - Some Characteristics of the Current System 10
   - Modeling Issues and Assumptions 14

3. UNCERTAINTY, SYSTEM PERFORMANCE, AND "LONG SUPPLY" 24
   - Uncertainty 24
   - The Role of Uncertainty in Long Supply 25
   - Implications of These Explorations for Management Strategies 31

4. UNCERTAINTY AND MANAGEMENT ADAPTATIONS 32
   - Structure of Requirements Evaluations 32
   - The Quick Response Option 34
   - Implications of the Evaluations for Spares Requirements 41

5. TOWARD IMPROVED ESTIMATION OF SPARES REQUIREMENTS 43
   - Enhancing the Responsiveness of Depot Repair 43

Designated Cannibalization 44
Improving Demand Modeling and Forecasting 46
Adding Engines to the Requirements Database 48
Data and Data Processing Issues 49
Developing a Configuration Database 49
Consolidating WRSK 51
SRU Safety Stock at the Depot 54

6. ESTIMATING DEPOT COMPONENT REPAIR REQUIREMENTS 55
Data and Data-Processing Issues 56
Applicability of Drive to Specific Shops and Workloads 58
An Analysis of DRIVE's Repair Requirements Estimation 58
DRIVE and the Critical Item Program 67

7. TRADEOFFS AMONG SPARES, REPAIR, AND TRANSPORTATION 69

8. CONCLUSIONS AND RECOMMENDATIONS 71
Recommendations 72

APPENDIX: ADDITIONAL DATA ON REPAIR REQUIREMENTS 77

BIBLIOGRAPHY 87

FIGURES

S.1. Effects of Uncertainty, Management Adaptations, and Prohibited WRSK Withdrawal on Peacetime Performance ix
S.2. Designated Cannibalization Pays Off xiii
2.1. The Inventory System Comprises Many Low-Demand Items 12
2.2. The Mix of Assets Recommended for Procurement by the Requirements System 13
2.3. A Schematic Model of the Current System 14
2.4. Scatter Diagram of VTMRs and Item Pipelines 21
2.5. Scatter Diagram of VTMRs and Item Pipelines (Enlarged) 21
3.1. A Notional Characterization of Budgetary Requirements Profiles Resulting from the Steady-State Assumption 25
3.2. Probability of Long Supply vs Demand Rate 27
3.3. Probability of Long Supply vs NRTS Rate 28
3.4. Probability of Long Supply vs VTMR 28
3.5. Probability of Long Supply vs Changing Mean Demand for a Medium Demand Item 29
3.6. Probability of Long Supply vs Changing Mean Demand for a High Demand Item 30
4.1. Schematic Representation of Evaluation Software System 33
4.2. Peacetime Performance with WRSK Withdrawals 35
4.3. Effects of Churn on System Performance with WRSK Withdrawals 36
4.4. Effects of Churn on Peacetime Performance, No WRSK Withdrawals Allowed 38
4.5. Estimated Wartime Performance, WRSK + Primary Operating Stock 39
4.6. Effects of Churn on Wartime Performance, WRSK + Primary Operating Stock 40
4.7. Performance of Alternative WRSK Withdrawal Policies 40
4.8. Effects of Churn on Performance of Alternative WRSK Withdrawal Policies 41
5.1. Designated Cannibalization Pays Off 46
5.2. Effects of Consistency in the Number of Users 51
5.3. Savings Achievable with Alternative WRSK Policies 53

6.1. Total IM, X21, and Negotiated Requirements vs the Keep-Up Requirement 60
6.2. Repair Requirements for the CI Test Station 61
6.3. Asset Position of the Fire Control/Navigation Panel (3533) 62
6.4. Asset Position of the 3160 Panel 62
6.5. Asset Position of the Inertial Navigation Unit (6645) ... 64
6.6. Estimation of the Catch-Up Requirement 66
A.1. Repair Requirements for the RF Test Station 77
A.2. Asset Position of the 63BG0 Indicator (RF1313) 78
A.3. Asset Position of the Radar Antenna (RF4630) 79
A.4. Asset Position of the LPRFs (RF296X) 79
A.5. Repair Requirements for the DI Test Station 80
A.6. Asset Position of the Radar Control Panel (DI6872) 81
A.7. Repair Requirements for the PP Test Station 82
A.8. Asset Position of the Pneumatic Sensor (PP6398 or 7835) 82
A.9. Asset Position of the Central Air Data Computer (PP1859) 83
A.10. Asset Position of the Central Interface Unit (PP4855)... 84
A.11. Asset Position of the Missile Release Interface Unit (PP1499) 85
A.12. Asset Position of the Stores Control Panel (PP6879) 85

TABLES

S.1. Performance of Improved Techniques Compared to Current System vii
S.2. The Quick Response Option vii
2.1. Numbers of Users, F-16-Peculiar Items 18
4.1. The Quick Response Option 34
5.1. Percentage Improvement in RMSE and MAD of Improved Techniques over Current System 47
5.2. Performance of Improved Techniques Compared to Current System 48
6.1. Comparison of the TAC Critical Item List with a DRIVE List 68
7.1. Distribution of Depot Repair Times 70

GLOSSARY

AFLC. Air Force Logistics Command, now the Air Force Materiel Command.

AFMC. Air Force Materiel Command, formerly the Air Force Logistics Command.

Aircraft availability. An aircraft is defined to be available if it is not waiting for a component to be repaired or be shipped to it.

Aircraft Availability Model (AAM). The system of software imbedded in D041 that is used to compute requirements for safety stock of selected recoverable items.

ALC. Air Logistics Center. One of five subcommands of AFMC.

Anticipated stockage posture. The asset position that is anticipated to eventuate in the inventory system an average lead time beyond the buy point.

AWP. Awaiting parts, the status of a component whose repair cannot be completed until repair parts are available.

BCS. Bench check serviceable, the status of a component removed from an aircraft in the belief that it is unserviceable but subsequently no defect is found in its condition or operation.

Beddown. A term denoting the allocation of aircraft or other weapon systems by type to locations.

Bias. The property of a statistical estimator such that its mathematical expectation differs from the numerical value of the parameter it is used to estimate.

BP15. Budget Program 15, a category of appropriated funds allocated to recoverable aircraft replenishment spares.

Churn. The sum total of all changes in the requirements database from one point in time to the next: the appearance of new items, the disappearance of items formerly in the database, and changes in item characteristics, e.g., demand rates, repair times, and NRTS rates.

Compound Poisson process. A stochastic process in which the numbers of arrivals that occur in disjoint time intervals of equal length are described by the Poisson probability distribution, and

the number of events that occur with each arrival is described by a separate, usually different, probability distribution. The number of events that occur with each arrival is called the *compounding random variable.*

Consumable. The property of a part or material such that it is discarded after failure or is consumed in use.

CONUS. The Continental United States.

CSIS. The Central Secondary Item Stratification, AFMC's system that is incorporated in D041, along with the Aircraft Availability Model, for computing requirements for recoverable spares.

D028. AFMC's Central Stock Leveling System that allocates stock levels for recoverable items to the bases and the depot.

D041. AFMC's system for computing requirements for aircraft recoverable spares, formally entitled the *Recoverable Consumption Item Requirements System.*

D165. AFMC's Mission Capability Requisition Status Reporting System, which reflects shortages of components of aircraft and other items of equipment.

Demand. In the Air Force's Recoverable Consumption Item Requirements System, the system used to compute spares and repair requirements, *demands* are defined as removals of components from their next higher assemblies, but they exclude those components that are declared to be serviceable after subsequent bench check. They also exclude part removals to facilitate other maintenance, etc.

Depot turnaround time. The average or actual time between the removal of a component from an aircraft and its repair at the depot and return to serviceable stock at the base.

DRIVE. Distribution and Repair in Variable Environments, AFMC's system of software for prioritizing the repair of aircraft components and allocating the serviceable assets emerging from repair to the bases and the depot. It is also useful for computing component repair requirements.

K004. AFMC's database that contains program data, i.e., flying hours, aircraft allocations, etc.

Long supply. In this context, long supply is defined to be all assets on hand in the inventory system above the number determined to

be needed by the requirements computation a lead time beyond the buy point.

LRU. Line-replaceable unit, a part or assembly that is typically removed directly from an aircraft when undergoing maintenance other than adjustment, calibration, or servicing.

MAD. Mean absolute deviation. The average unsigned difference between a set of estimators and the true values of the parameters being estimated.

Master stock number. The stock number assigned to the preferred item in a set of two or more interchangeable items.

MDS. Model-design-series, the system used to designate aircraft.

MICAP. A term denoting the mission capability effect of a part shortage.

MISTR. Management of Items Subject to Repair, AFMC's system for managing the depot-level repair of assets.

Moving average. The statistic formed by the mean of a fixed number, n, of observations of a stochastic process where the most recent n observations are summed and divided by n. As observations accumulate over time, the latest observation is added to the sequence and the $n + 1$st observation, counting backwards, is discarded.

NATO. North Atlantic Treaty Organization.

Negative binomial probability distribution. A probability distribution of the discrete type that may apply to situations in which events occur at random but the variance of the numbers of events in nonoverlapping time intervals of equal length is higher than allowed by the Poisson distribution. Its density function is given by

$$\binom{r + x - 1}{x}(1 - p)^x p^r, \; x = 1, 2, \ldots$$

NFMC. Not fully mission capable, the status of an aircraft whose ability to perform its mission is degraded.

NMCS. Not mission capable, supply. The status of an aircraft that is incapable of performing its mission because of a shortage of parts.

NRTS. Not repairable this station. The designation given to a repairable part whose repair is beyond the capability of maintenance at a particular location.

NSN. National stock number, a unique number assigned to identify a particular kind of item. It comprises a four-digit federal stock class, a nine-digit national item identification number, and a two-character suffix denoting its materiel manager.

NSO item. A numerical-stockage-objective item, one whose demand rate is low and whose requirements are based on levels jointly established by the item manager and equipment specialist.

OIM. Organizational and intermediate maintenance.

OIMDR. Organizational and intermediate maintenance demand rate, which reflects the number of demands per unit of past installed item program over the previous eight quarters. For aircraft recoverable components, past installed program is typically expressed in hundreds of flying hours absorbed by installed items. A demand is the removal of a component in the belief that it is defective, although when the component is judged to be serviceable with no repair required by the intermediate repair activity, the demand is not counted.

OSD. The Office of the Secretary of Defense.

PAA. Primary authorized aircraft.

Pipeline. In the spares requirements context, the expected number of assets in resupply.

Poisson process. The most widely known and often used form of stochastic model with important mathematical properties that make it especially tractable and useful. It is described by the Poisson probability distribution whose density function is given by

$$\frac{e^{-\lambda}\lambda^{x}}{x!},\ x = 0, 1, 2, \ldots$$

POS. Primary operating stock, formerly known as *peacetime* operating stock.

Power function. A mathematical function of the form $y = ax^b$.

QPA. Quantity per application, the number of parts of a particular type installed on the part's next higher assembly.

REALM. The Readiness/Execution Availability Logistics Module, the software module of WSMIS that computes requirements for war readiness spares.

Recoverable parts. A class of parts subject to repair when they fail, rather than being discarded or consumed in use.

Resupply. The state of parts that are in base repair, depot repair, shipment from one location to another, or have been condemned and whose replacement is pending.

RMSE. Root mean squared error, computed as the square root of the average squared difference between a set of estimators and the true values of the parameters being estimated, a popular measure of predictive accuracy. Also referred to as RMSD, root mean squared deviation.

Safety stock. Spares authorized to accommodate the variability in the numbers of items in resupply.

SRU. Shop-replaceable unit, a subassembly of an LRU that is typically replaced during repair of the LRU.

Stationary process. A stochastic process whose parameters are invariant.

Stock level. Serviceables on hand plus due-ins minus due-outs.

VSL. Variable safety level, the Air Force's implementation of a method for estimating spares requirements developed by C. C. Sherbrooke of RAND called METRIC (Multi-Echelon Technique for Recoverable Item Control).

VTMR. Variance-to-mean ratio, defined as the unbiased estimator of the variance of a process divided by its mean.

WRSK. War readiness spares kit, an airlift-deployable set of spares to support squadrons deployed in contingencies.

WSMIS. AFMC's Weapon System Management Information System.

ZI. The zone of the interior, i.e., the 48 contiguous states.

1. INTRODUCTION

The principal aim of the research described in this report is to understand how demand uncertainty and the effectiveness of management adaptations in coping with it should affect the estimation of requirements for aircraft recoverable spares and depot repair in the Air Force.[1] This work is limited in scope to aircraft replenishment spares (in contrast to components of ground radars, say, or to initial spares), although some illustrations are included that encompass all recoverable items.

The issue is a major one. In the mid-1980s, the Air Force spent roughly $5 billion dollars annually on purchases of aircraft recoverable spare parts and their depot-level repair. About $3 billion of this budget was for procurement and about $2 billion for repair.[2] At the time of this writing, the Air Force is trying to cope with dramatic reductions in its budgetary resources.

The assumptions underlying the Air Force's estimation of its budgetary requirements for recoverable spares are inconsistent in important ways with the realities of actual logistics management practices. For example, in estimating spares requirements, it is assumed that parts shortages are randomly distributed among aircraft to which they are applicable; but in practice the shortages of many parts are consolidated into as few aircraft as practicable, i.e., parts are *cannibalized*. In actual practice logisticians adapt to circumstances inducing shortages, not only by cannibalization, but in other ways: expediting transportation and handling, expediting repair, withdrawing assets from war readiness spares kits (WRSK), polling other bases to find a serviceable spare asset and invoking *lateral supply,* and even drawing upon the resources of manufacturers in some urgent circumstances. Such adaptations are not reflected in the requirements system's model of the logistics system nor in its underlying assumptions; therefore, the logistics system delivers levels of performance that typically substantially exceed the levels specified in the requirements computation.

[1]A recoverable asset is one that is subject to repair when it fails, in contrast to a consumable, which is discarded upon failure or consumed in use.

[2]These numerical values were derived from our replication of AFLC's requirements computation done in the third quarter of fiscal year 1986.

The assumption is also made that future demand rates will be the same as past demand rates, future pipeline times the same as past times, and so on, i.e., the logistics system is viewed in some ways as a steady-state system. The assumption of a steady-state world extends implicitly to the requirements database. In the absence of explicit interventions by the item manager (IM) or equipment management specialist (EMS), whatever demand rates, NRTS (not-repairable-this-station) rates, repair times, and other important item characteristics are reflected in the current database are assumed to be applicable to the item at a future time. From one year to the next, however, our estimates of item characteristics change in ways that are consistent with changing, evolving Air Force experience. Moreover, many items that are in the database now will not be there a year from now; conversely, new items will appear in the database a year from now that are not there this year. We describe the sum of all such changes in the database from one year to the next as *churn*. Churn is an important factor in requirements estimation, so important that any method that ignores it is vulnerable to large predictive errors. Churn acts in the opposite direction in its effects on system performance than management adaptations do. It results in substantially poorer performance than a steady-state assumption suggests. Moreover, it contributes to the problem of *long supply*, i.e., an overabundance of certain assets in the inventory system.

Unfortunately, improved modeling and forecasting of spare parts demands can be of only limited help in coping with uncertainty, although improvements are clearly worth pursuing. Improved modeling and forecasting techniques are discussed at length in a companion report (Adams et al., 1993).

Although the understatement of our uncertainty about the future and the understatement of the system's ability to adapt to urgent, unanticipated demands tend to be offsetting, as we will show, the pooled error tends to be conservative, i.e., the system can reasonably be expected to do considerably better than the goals specified to the requirements computation. It is important to point out in this context that the specification of an aircraft availability goal of, say, 83 percent for the F-16 aircraft, doesn't necessarily mean that the Air Force would be satisfied with 17 percent of its F-16 aircraft unavailable for lack of recoverable components. Clearly, if one repeatedly specifies a goal of 83 percent in the face of an observed performance level of, say, 96 percent, one is implicitly specifying a 96 percent goal for which the number 83 is merely a numerical transformation to account for lack of fidelity in the computational model.

The requirements computational process, from the specification of goals to its numerical computations and modeling, is fraught with approximations and assumptions to the extent that one cannot accurately estimate the performance that will eventuate from any particular mix of spares in the inventory system without fairly sophisticated evaluation devices external to the requirements computation itself. This fact shaped our approach to this research for which we built a system of software that replicates the Air Force Materiel Command's (AFMC's) requirements computation including its central stock leveling system (D028). This enables us to estimate the stockage posture that will eventuate from any specified mix of spares procurements. In this work we used the F-16 aircraft as a case study, but we replicated the requirements computation for the entire recoverable spares inventory system to account for common items and estimate the total requirement.

There is an important missing link in our ability to estimate future stockage postures. We cannot account for the behavior of item managers; therefore, we cannot model how the "shopping list," i.e., the list of recommended spares procurements that emerges from the requirements computation, gets translated into purchase requests by item managers and others. Our concern is with enhancing the rationality of the requirements computation itself. Human intrusion in the process may act either to improve the performance of the mix of spares purchased or to diminish it; we simply don't know, although we suspect that some of both occurs. It must suffice to say here that we focus our attention on getting the computation right.

Data and data-processing issues also have important implications for the quality of the results yielded by the requirements computation. Some of these issues are discussed at length in a companion report (Abell and Finnegan, forthcoming).

STOCK FUNDING OF RECOVERABLES

It seems unlikely that recoverable spares requirements estimation will be affected significantly by the Air Force's implementation of stock funding of recoverables. In past years, there was a sense of urgency to get funds committed early in the fiscal year. This is likely to change because of the revolving nature of the stock fund, but there does not seem to be any characteristic implicit in stock funding that would change the fundamental logic of the spares and repair requirements estimation system.

FORCE DRAWDOWN

The number of active aircraft in the Air Force is anticipated to decline in the years ahead. Obviously, this decline will cause many assets to become excess to current requirements, and spares and repair requirements to decline over some period. Anticipated changes in force size have no bearing on this work. We have tried through this research to enhance the rationality of the process. Although the scale of things is bound to change, the logic of the requirements estimation problem will not. Thus, what we recommend here is as valid for a small force as it is for a larger one.

WHAT FOLLOWS

In this report, we explore the current spares and repair requirements system, its important characteristics, its assumptions and their implications, and alternative policies and approaches that we believe can achieve desired levels of system performance at reduced cost, while mitigating the effects of some of the specific problems that are so troublesome to the Air Force in the austere budgetary environment it now faces. In Section 2 we describe the essential character of the spares requirements estimation problem and the current requirements system in the context of the Air Force's recoverable spares management system, and how the difficulty of the problem is compounded by the formidable levels of uncertainty underlying the demands for aircraft recoverable spare parts and their depot-level repair. In Section 3 we discuss the role of uncertainty in spares requirements estimation and its effects on system performance.

In Section 4 we describe some evaluations of management adaptations such as lateral supply, cannibalization, WRSK withdrawals, and expedited handling and transportation of assets. We describe an improved approach to estimating spares requirements in Section 5 and offer several specific recommendations for modifications to the current approach.

Sections 2, 3, and 4 deal primarily with the estimation of spares procurement requirements. Component repair requirements are also of interest here because, with the advent of stock funding of recoverables in the Air Force, spares procurement and repair requirements will be funded from the same source. In Section 6 we discuss an improved approach to estimating depot component repair requirements and describe the specific tasks that need to be undertaken to implement it. The opportunities for tradeoffs among spares, repair, and transportation are important to any effort to achieve specified levels

of system performance at least cost; the issues related to such trade-offs are discussed in Section 7. Section 8 contains our conclusions and recommendations.

2. A BRIEF OVERVIEW OF SPARES REQUIREMENTS ESTIMATION

We begin with a brief, somewhat simplified overview of the recoverable spares management system. At bases that possess aircraft, discrepancies arise in aircraft maintenance and operations that lead to the diagnosis of problems with the hardware components of the aircraft. Sometimes those problems are diagnosed and repaired on the aircraft with an adjustment, minor repair, or replacement of consumable parts. Other times they are corrected through the replacement of a recoverable component, typically described as a line-replaceable unit (LRU). When an LRU is removed from an aircraft in the belief that it is defective, a base *repairable generation* (or *rep gen*) is said to occur. In the Air Force's traditional maintenance structure, the repairable[1] component is delivered to the intermediate-level maintenance activity for fault diagnosis and repair. If the repair is beyond the capability of intermediate level maintenance, it is declared to be NRTS and is usually returned to the depot-level repair activity for repair and return to serviceable stock.

SOME COMMON TERMS

Unserviceable assets en route from the base to the depot are said to be in *retrograde*. After repair and return to serviceable condition, assets are available for issue. A base submits a requisition to the depot for a replacement asset when it sends a NRTS asset back to a depot repair facility. Assets on order by the base are called *due-ins*. Assets that have been requested from base supply by base maintenance but not provided by supply are called *due-outs* or *base backorders*. The base's generation of replenishment requisitions is governed in part by a *reorder policy*, which in the case of recoverable items is described as a *continuous review* reorder policy with an *order quantity* of one. This implies that whenever the base's stock on hand plus due-ins minus due-outs falls below the *stock level*, a replenishment requisition is sent to the depot by the base. If the depot fills the requisition upon receipt, the transaction is referred to as an *issue*. If it does not, a *depot backorder*, or *depot due-out*, is created. The allocation of a serviceable asset to a depot backorder is called a *backorder release*. Note

[1]The term *repairable* is defined to mean defective, i.e., unserviceable, in contrast to *reparable*, an awkward synonym for *recoverable* popularized by the other services.

that depot backorders do not necessarily imply that an aircraft is short of a component; it simply means that the depot owes the base an asset. On the other hand, a base backorder does imply that a component is missing from an aircraft. These shortages are sometimes referred to as *holes* in aircraft.

THE CONCEPT OF ITEM PIPELINES AND RESUPPLY

The models of the inventory system used in computing spares requirements are built around the concept of an item's *pipeline*. The notion of a pipeline in this context is different from the generic usage of the term to describe a conduit. Suppose that one component of a particular kind of aircraft, say an F-16A, is a radar transmitter. Suppose, too, that there is one radar transmitter installed in each F-16A aircraft. When a transmitter fails, or is suspected of failure, it is generally removed from the aircraft, in the manner we have already described, for repair in the intermediate-level repair activity (typically located on the same base). It is either repaired there or sent back to the depot (or a contractor facility) for repair. Eventually it emerges from repair and is typically allocated to some base as a replacement asset. If there were no serviceable spare transmitters in the supply system, the base at which the failure and removal occurred would have to wait for the repair of the actual component removed from the aircraft. To avoid such delays we provision the system with *spares*, transmitters to be installed in aircraft to replace defective ones while waiting for them to be repaired. Thus some number of transmitters may be in *resupply*, i.e., in intermediate-level repair, in retrograde to the depot, undergoing depot repair, in shipment from the depot to a base, or in the condemnation pipeline. The expected number of items of a particular type in resupply is commonly referred to as the *pipeline*, or *item pipeline*. It is simply a mathematical expectation of the number of assets, transmitters say, that would be missing from aircraft if there were no spares in the supply system (assuming the flying hour program was actually flown).

The pipeline is the basis for the computation of spares requirements. In practice, the number of assets actually in resupply varies over time. If the number of transmitters removed from aircraft every day never varied, and their repair and return to the base always took exactly the same amount of time, the pipeline would have little or no variability, and the system could operate effectively simply by buying enough spares to fill the pipeline.

In actuality, however, there is great uncertainty in every segment of the pipeline, and in all of the factors used to estimate its numerical

value. Demand rates, NRTS rates, repair times, shipment times, and other measures of system performance change over time. Depot repair times and the times that elapse between an asset's arrival at the depot and its induction into the depot repair shop also vary greatly among shops and items and over time. Thus, although the mathematical expression for the expected number of assets of a given type in resupply is simple and straightforward, it is very difficult to "get it right" in a practical sense; that is, the world often eventuates in ways that we cannot predict. When resupply times are longer than forecast, spares shortages may develop. When repair activities are able to respond more quickly than planned, more serviceables may be available than anticipated. The vagaries of the processes underlying the richness or paucity of serviceable spares in the supply system are substantial. Our uncertainties about factors that affect the numbers of assets in resupply greatly complicate the problem of deciding on the appropriate levels of investment in aircraft recoverable spares, and on the best mix of spares to buy to achieve our aircraft availability goals at least cost.

Because of our uncertainties about the actual number of assets in resupply, we invest in *safety stock* as well as in the assets required to fill the pipeline. Thus safety stock is intended to protect the inventory system against *variability* in the number of items in resupply. Safety stock also helps protect the system against errors in estimating item pipelines, another important source of uncertainty.

The problem, though simple enough to state, is difficult to solve: *Determine the level of investment in aircraft recoverable spares and the best mix of spares to buy to achieve specified aircraft availability goals at least cost.* This is the fundamental problem addressed in the research reported here. It is made difficult by its sheer dimensions, the quality of data used to support requirements computations, and the complexity of the logistics system whose behavior the requirements system tries to model, but most of all, perhaps, by the formidable uncertainties that underlie our estimates of the factors that shape the pipelines of aircraft recoverable spares.

The approach AFMC uses to estimate procurement and repair requirements for these assets has long absorbed the attention of many persons in the logistics research community. Progress has been made on several issues, most notably, perhaps, on the logic underlying the computation of safety stock for an important subset of items in the recoverable inventory system. Despite this progress and, ironically, in part because of it, the Air Force has been the object of heavy and persistent criticism, especially in recent years, owing to a situation typi-

cally described as *the long supply problem* that people often perceive as symptomatic of overstated requirements. As we will attempt to show, such criticism is only partly justified for reasons that we will describe in Section 3. We also discuss the long supply problem, its implications, and its underlying causes at greater length in Section 3.

THE ROLE OF UNCERTAINTY IN SPARES REQUIREMENTS ESTIMATION

In his insightful discussion of uncertainties and a taxonomy of approaches to coping with them, Hodges (1990) defines two principal types of uncertainties:

- Statistical uncertainty, defined as variability observed in repeatable phenomena, and
- State-of-the-world uncertainty, defined as uncertainty about phenomena that are not repeatable, not observed or observable, or both.

In many applications in the decision sciences, in analysis, in system design, etc., we find solutions to problems involving statistical uncertainty. Such is the case in current spares requirements computational models. The real problem, however, is plagued by the more imposing state-of-the-world uncertainty. While it is difficult enough to cope with statistical uncertainty, it is far more difficult to cope with state-of-the-world uncertainty. Often the sources of uncertainties are exogenous to the spares requirements system: budgets, flying hour programs, or weapon system acquisition programs are changed; unanticipated, urgent demands are generated by an unforeseen and unplanned for event such as a military contingency; individual aircraft components sometimes experience dramatic and sudden increases in their failure rates; repair capability may suddenly be lost temporarily, resulting in sharp increases in the number of items in resupply; transportation resources may be diverted for a time to higher priority needs; and so on. It is easy to offer many examples of events that degrade the performance of the inventory system and its ability to provide planned levels of logistics support to the combat force. We are only very rarely able to anticipate such events. The only forecast one may make confidently is that the future will almost certainly eventuate differently than we anticipate.

It is important, then, to make decisions about resource allocations and system design that enhance the *robustness* of system performance, i.e., the ability of the system to deliver adequate performance

in a variety of possible futures, rather than basing those decisions only on the expected performance of the system in a prescribed planning scenario. It is safe to say that uncertainty is the most dominant factor compounding the difficulty of estimating spares requirements accurately. The formulation of the computational models in the spares requirements system account, in part, for uncertainties of the statistical variety. Some provision is made for human intervention to accommodate state-of-the-world uncertainties, but many such uncertainties are difficult even to foresee, much less quantify. A major thrust of this work has been to quantify the effects of uncertainty on the relationship between investments in aircraft recoverable spares and logistics system performance measured in aircraft availability, and to account for some management actions that help the system overcome those effects.

SOME CHARACTERISTICS OF THE CURRENT SYSTEM

In a review of the March 1986 and 1989 D041 databases,[2] we tried to characterize the system statistically to provide the reader with some intuition about the size of the system and other important characteristics.

The Air Force's recoverable item inventory system is large. The 1986 database contained 178,749 item records. Each represents a unique item that is assigned a national stock number (NSN). The NSNs that appear in the database are at least subgroup master items, i.e., they are typically the most preferred, latest configurations of a group of one or more NSNs that have the same form, fit, and function. For example, the radar transmitter we discussed previously might have an earlier configuration that still does the job, so to speak, but isn't quite the latest, preferred configuration. The earlier version may not explicitly appear in the database. There are roughly 35,000 additional NSNs that are not in the database but are interchangeable with or substitutable for NSNs that are in the database. Data reflecting demands, NRTS actions, repairs, etc., for these items are pooled with the data of the included items.

A surprising characteristic to many who are unfamiliar with the inner workings of recoverable item management is the fact that, of these 178,749 NSNs, the records of 131,283 (73.4 percent or almost three-fourths) of them reflect an organizational and intermediate

[2]D041 is the data system designator for the *Recoverable Consumption Item Requirements System*. It processes and stores inputs from many other AFLC standard data systems to support the computation of recoverable spares requirements.

maintenance (OIM) demand rate[3] of 0.0. Items without a demand in the past eight quarters are said to be *inactive* items. Thus a fairly small proportion of the items in the inventory system absorb most of the demands.

The demand rate also affects the categorization of items as insurance items, numerical-stockage-objective (NSO) items, deferred-disposal or retention items, and contingency-retention items. An insurance item is one that is not subject to wear and tear but may be required as a result of an accident, fire, crash, or other unforeseen cause. Their requirements are not computed because they typically have zero demand rates; rather, their requirements are based on token levels (generally three or less).

An NSO item is simply one with a low demand rate; its requirements are based on levels jointly established by the item manager and equipment specialist.

A deferred-disposal or retention item is one part or all of whose computed excess quantities have been specified for retention. Contingency items are items with no future programs. Contingency-retention items are contingency items all or part of whose assets are being retained at the direction of higher authority for some purpose other than the original intended use. They are usually obsolete or nonstandard items.[4] All of these categories of items have either very low or no past demand. Figure 2.1 shows the distribution of items in several categories. Of the 106,473 BP-15 items[5] in the March 1986 D041 database, 59,285 (55.68 percent) of them were in these categories. The point of the chart is that requirements for less than half the items in the system are demand-driven.

Another important characteristic of the inventory system lies in the distribution of the dollar value (based on unit prices reflected in the database) of various categories of assets in the mix of spares bought

[3]The organizational and intermediate maintenance demand rate (OIMDR) reflects the number of demands per unit of past installed item program over the previous eight quarters. For aircraft recoverable components, past installed program is typically expressed in hundreds of flying hours absorbed by installed items. A demand is the removal of a component in the belief that it is defective, although when the component is judged to be serviceable with no repair required by the intermediate repair activity, the demand is not counted.

[4]These definitions are discussed at greater length in a training document, Air Force Logistics Command (1984).

[5]BP-15, or budget program 15, is associated with aircraft replenishment spares, in contrast to the initial spares that are purchased early in the life of a weapon system.

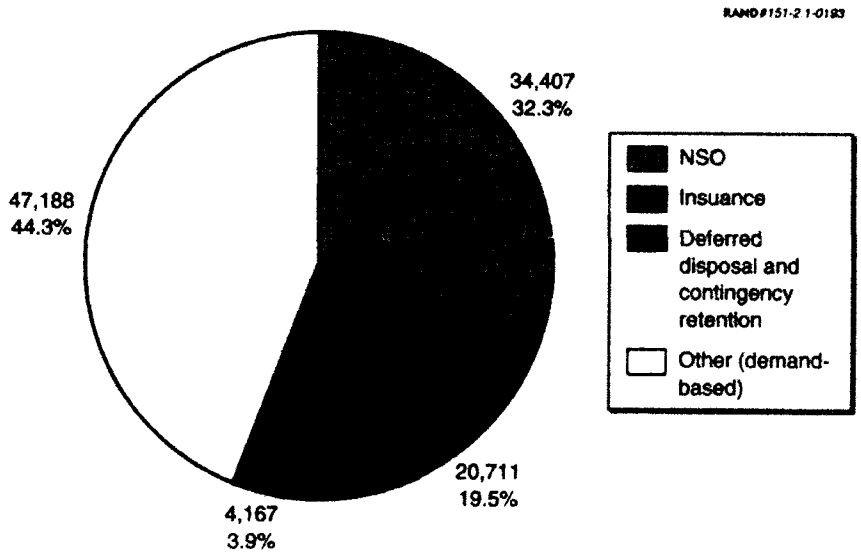

Figure 2.1—The Inventory System Comprises Many Low-Demand Items

in contrast to the mix of spares already in the inventory system. To categorize the mix of spares bought, we divide assets into categories of pipeline, safety stock, and negotiated quantities. To categorize the mix of spares in the inventory system, we pool safety stock with items in "long supply," i.e., items that are in excess of the requirements computed a lead time beyond the buy point, ignoring any possible longer-term asset applications.

Figure 2.2 shows the distribution of dollar value among the categories of items recommended for procurement by the requirements computation done with the March 1989 D041 database. The recommended buy is split roughly three-fourths to fill pipelines and one-fourth for safety stock and negotiated levels.

The mix of assets in the inventory system is quite another matter. Much of what is on hand in the inventory system is not likely ever to be demanded. In the view of the requirements computation just described, roughly three-fourths of the assets are not judged to be needed a lead time beyond the buy point. The difficulty, obviously, is

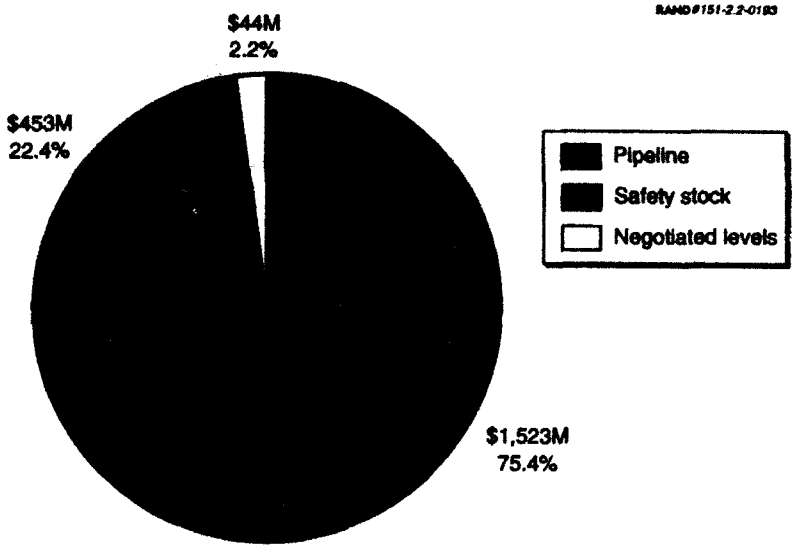

Figure 2.2—The Mix of Assets Recommended for Procurement by the Requirements System

that we don't know which assets will be demanded and which won't be. This situation is consistent with what we know about churn. Recall that churn is simply the sum of all the changes in the requirements database from one year to the next. Thus, when we compute the requirement at any specific time, we have, in a sense, forecast the future. Because significant time passes between the computation and the delivery of assets from contractors, the mix of spares we bought is almost invariably the wrong mix at the time they are delivered. Our estimates of item characteristics, the force beddown, flying hour programs, modifications to items, and many other factors that shape the desired mix of spares evolve and change. Thus, many spares that we bought in the past, especially the distant past, are increasingly less likely to be relevant to our current needs as time passes, and many of them will be insufficient in quantity because of increased failure rates and other factors.

The important points of this discussion are:

- The mix of spares we think we will need at any point in time in the future changes as time passes, our experience evolves, and that experience is incorporated in the requirements database (Lippiatt, 1991).

- The mix of spares in the inventory system reflects procurements over many past years and cannot be expected to reflect our current needs.

Gains may be possible in mitigating the effects of uncertainty on long supply and, in fact, we explore several such initiatives later in this report.

MODELING ISSUES AND ASSUMPTIONS

Figure 2.3 illustrates the workings of the Air Force's current system for computing spares and repair requirements. Assets and requirements are partitioned into two classes: primary operating stock (POS) and war readiness spares. A subset of WRM that comprises airlift-deployable sets of spares called war readiness spares kits (WRSK) are of special interest in the system. The computation that is fundamental to estimating POS spares and repair requirements is called the Central Secondary Item Stratification (CSIS). It provides budget summaries and repair summaries on a quarterly basis. The CSIS is augmented by the aircraft availability model (AAM) that computes the requirement for POS safety stock.

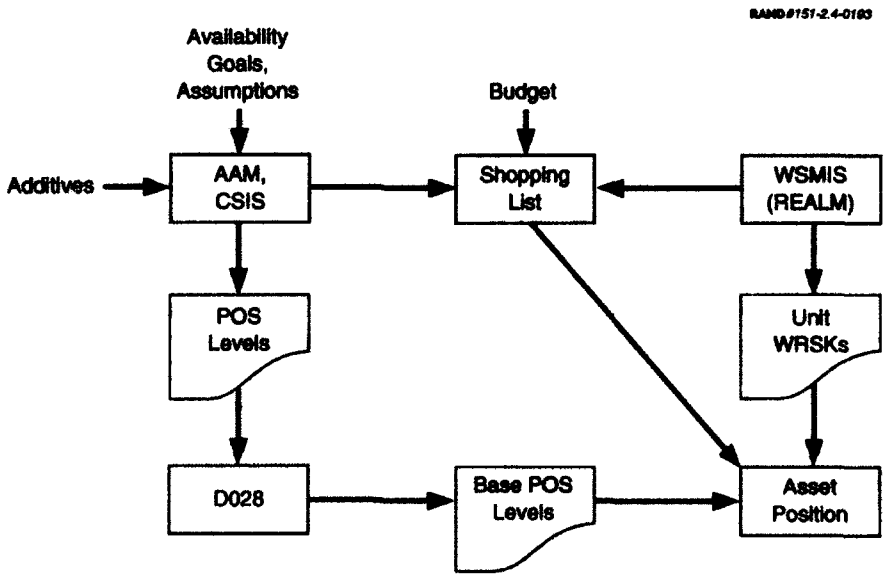

Figure 2.3—A Schematic Model of the Current System

The required number of assets of each type that is computed by this system is fed to the AFMC central stock leveling system (D028), which allocates stock levels to the bases and the depot. WRSK requirements are computed in the Weapon System Management Information System (WSMIS) and are adjusted through negotiations with the major commands. A base's future asset position comprises both POS and WRSK assets and is shaped largely by these two requirements estimation processes. At the time of the computation, we refer to this asset position, or stockage posture, as the *anticipated* asset position. It eventuates, on average, roughly three years after the requirements computation takes place.

Safety stock is an additive to the CSIS requirement that is now computed by the aircraft availability model (AAM). Before 1975, AFLC used a fixed-safety-level computation, i.e., each item received essentially the same level of protection against variability in demand. In that year, AFLC implemented a new approach to these computations that was called *variable safety level* (VSL). As the name suggests, individual items did not all receive the same level of protection; rather, the amount of safety stock computed was greater for low-cost items than for high-cost items. This *system approach* to computing safety stock was intended to deliver better system performance per dollar invested than the fixed-safety-level approach. Conversely, it would compute a lower-cost mix of spares for any specified level of system performance.

The objective function used in the VSL computation was total expected base-level backorders. The VSL approach was AFLC's adaptation of a model called METRIC (Multi-Echelon Technique for Recoverable Item Control), developed by Craig C. Sherbrooke, formerly of RAND, in the mid-1960s (Sherbrooke, 1966).

AFLC moved to the use of the AAM in recent years. It was developed by the Logistics Management Institute and is used for computing safety levels for selected items. The AAM extends the expected-backorder logic to an expected aircraft availability objective function (O'Malley, 1983).

There are several important technical issues in the way the current system estimates spares and repair requirements. In any system as complex as the Air Force logistics system, assumptions and approximations are virtually always needed in models of the system. The requirements system is no exception. We will explicate some of those assumptions and approximations here to (a) provide insights about approximations and assumptions in the computations and the likely directions of their differences with observed system performance, and

(b) develop the logical basis for suggestions about how the system might be improved. In the several pages that follow, we will point out that the current system's specification of probability distributions of the numbers of assets in resupply is, on balance, conservative from the standpoint of logistics support; i.e., it computes requirements for more spares than needed to achieve specified aircraft availability goals because, by design, it does not account for several of the management adaptations that are routinely practiced even in peacetime to help mitigate the effects of demand variability and erroneous forecasts. As we will show, those management adaptations have powerful effects on performance.

Churn also has important implications for spares requirements, and, to a lesser extent, for repair requirements estimations, that are not now accounted for. It also has profound implications for outyear requirements estimations as well as the phenomenon of long supply.

The spares and repair requirements system assumes that transportation times will be the same as they have been historically; i.e., historical data determine the times that are specified to the requirements computation. While certainly not an unreasonable assumption for normal peacetime operations, the computational system enables us to specify alternative levels of transportation system performance that could be tied to alternative configurations of, and resource allocations to, the transportation system, depot repair, and materiel support to depot repair, thus enabling decisionmakers to explore alternative resource allocations among spares, repair, and transportation. This ability could pay handsome dividends in total logistics system performance and cost. It suggests viewing transportation and repair system performance as decision variables rather than as constraints.

Data Problems

Some components that are removed from aircraft in the belief that they are defective are subsequently judged by the maintenance activity to be serviceable after some fault-diagnosis procedure or functional check. As a matter of policy, the Air Force has decided not to provide spares to protect this portion of the base repair pipeline; therefore, these actions are ignored in the estimation of OIM demand rates. This deliberate departure from reality causes the system to buy fewer spares for items with high BCS rates.[6] As will be shown in

[6]Bench-check-serviceable rates, i.e., the proportion of component removals for cause that turn out to be judged serviceable.

Section 4, the Air Force, on balance, overinvests in spares; therefore, one might conclude, it would be imprudent to include BCS demands in the base repair pipeline because it would induce even higher levels of investment. The problem is that ignoring BCS demands affects items differentially because BCS rates vary across items. Thus a wiser policy might include BCS demands and rely on other approaches to constraining investment levels. In this way, a more cost-effective mix of spares might be achieved.

We observed substantial differences in item programs between those reflected in D041 and those computed using K004 data[7] and the application file.[8] We found that 25 percent of F-16 item programs differed from the computed values by 33 percent or more. Clearly, this issue needs priority attention.

In the logic used to construct the indenture file, whenever an SRU appears in the application file and the program is unable to find its next higher assembly it promotes the SRU to an LRU. The result is overinvestment in these SRUs because shortages of them are now assumed to affect aircraft availability directly. Notable among such SRUs are aircraft engine components, a large subset of items in the database, owing to the deliberate exclusion of records of engines and engine modules. This omission has an important effect on the computed spares mix. In fact engine components represent 21.4 percent of the cost of F-16 recoverable OIM pipelines.

The Partitioning of Item Pipelines Among Users

An important variable in the computation of safety levels is the number of locations over which an item's pipeline is allocated. The AAM makes an assumption for computational efficiency that is popularly called the *average base assumption;* i.e., it assumes n equal bases, where n is used to partition the item pipeline into equal parts. In the current system, n is set equal to the numeric value of a data element in the item's D041 "01" record called the *number of users.* The number of users for an item is defined to be the number of locations (bases) that experienced two or more demands for the item over the past 12 months. Thus its numeric value is a matter of chance. The most important characteristic of this number is the fact that *it does not represent the number of locations at which an item is exposed to*

[7]K004 is AFLC's system from which flying hour programs are extracted to support the requirements computation.

[8]The application file reflects all applications of an item to its next-higher assemblies in each model-design-series (MDS) aircraft to which it applies.

OIM demand. It typically *understates* the number of those locations. Moreover, in computing future safety-level requirements, changes in the force beddown may be inconsistent with the number of users reflected in the database, creating a logical disconnect.

It is simple to determine the number of bases at which an item is exposed to OIM demand if the item has an application percentage of 100 for all of its applications. The force beddown would yield such an estimate in a straightforward manner. If the item's application percentage is less than 100, however, it is not possible to determine the true number of locations because the requirements system does not know on which subset of aircraft the item is installed, nor does it know where the aircraft are.

Table 2.1 reflects the number of users assumed by the AAM for the set of items from the March 1989 D041 database that are peculiar to the F-16A/B, F-16C/D, or both. Each item had one or more demands in the past two years, are LRUs, and have application percentages of 100. There were 585 master stock numbers in this set. Note the rather remarkable understatement of the actual number of bases at which these items were exposed to OIM demand by the number of users in the D041 database. Even their maximal values were *invariably* smaller than the number of bases.

The implications of this phenomenon are not entirely clear; however, we did estimate the effects of using the number of bases as the number of users in a requirements computation in which we specified the usual aircraft availability goals. The result was an increase in the F-16 spares budget of 23.2 percent. In this analysis, 2.2 percent of the safety levels decreased, and 83.3 percent remained unchanged. Moreover, 14.5 percent of the safety levels increased, 6.7 percent by

Table 2.1

Numbers of Users, F-16-Peculiar Items

Aircraft MDS	No. of Bases, 1992	Number of Users		
		Avg	Min	Max
F-16A/B	31	4.6	1	26
F-16C/D	24	3.3	1	19
Both	48	9.4	1	35

more than 50 percent. However, when we constrained the budget to its original value, expected aircraft availability increased by 9 percent of the possessed aircraft.[9] We then tried both mixes of spares at the higher budget level and again found that aircraft availability using the number of bases as the number of users was higher by 9 percent of the possessed aircraft. One way in which the effects of lateral supply can be modeled is to reduce the number of users to represent the number of groups of bases worldwide within which lateral supply is routinely practiced. Understating the number of users as the current system does acts in the same direction and may roughly approximate such an approach, but one could easily make a persuasive argument for specifying the number of users "correctly" (that is, as the number of locations where an item is exposed to OIM demand), and modeling the effects of lateral supply in a way that is more consistent with reality, i.e., more consistent with the groupings of bases among which lateral supply transactions typically occur.

Unfortunately, we are not able to assess the effects of this idea simply because we don't have base-specific application percentages for items. Base-specific program data would also be needed. To allocate assets among bases to achieve specified aircraft availability goals, such data are needed anyway. If they were available, the question of how best to partition item pipelines could be resolved through analysis. We discuss this problem at greater length in Section 5.

Demand Modeling

The current system uses a negative binomial probability function to describe the distribution of the number of assets of each type in resupply, i.e., in the resupply pipeline. The distribution is described by its mean and variance-to-mean ratio (VTMR), although other parameterizations are possible. The VTMR is specified as a power function of the expected item pipeline which, in turn, is based on a mean demand rate that is an eight-quarter moving average of past demand. It has the form VTMR = ax^b where x represents the item's eight-quarter mean demand rate. The VTMR thus established is then used to describe the variance of a negative binomial probability function whose mean is the estimated mean of the item's pipeline, i.e., it is the expected number of assets of a particular type in resupply, and whose variance is computed from the specified VTMR. At the time of this writing, the values of a and b are 1.132477 and

[9]The aircraft availabilities were estimated using Dyna-METRIC Version 6. This approach is discussed at greater length in Section 4.

0.3407513, respectively. The computed VTMR is constrained to be between 1.01 and 5.0.[10]

There are several problems with this procedure. The observed VTMR is only an estimator of the true VTMR underlying the demand process (although this very statement implies a steady-state process, which is typically not the case). The observed VTMR is biased; the bias is always negative; it increases as a function of the true VTMR and decreases as the mean demand rate increases. The VTMR estimator itself has very high variance; thus it is not a very reliable estimator. No account is taken of the fact that the observed VTMR is an estimator. It is treated as a population parameter. The power function that is used to specify a VTMR in the requirements estimation process is, in fact, partly shaped by the estimator's bias function.

This modeling procedure does not fit the observed data very well. The steady-state assumptions made by the current procedure are also inconsistent with the true volatility in the process. No account is taken, for example, of the fact that the more coarsely the observed data are partitioned, the higher the estimate of the VTMR, and the higher the variance of the estimator, phenomena that cannot be accounted for by the steady-state view of the current system.

Figure 2.4 plots the observed VTMRs of quarterly demands experienced by 129 items stocked in the F-15C/D war readiness spares kit for 24 PAA squadrons against the items' expected pipelines. It also shows the power function used in the current requirements system to assign VTMRs to item pipelines. The power function bears very little resemblance to the actual data. Unfortunately, some procedure to dampen out the wild variability in the VTMR estimator is needed. We have explored alternative models of the demand process that we believe would improve the cost-effectiveness of the requirements computation. They are discussed in a companion report (Adams et al., 1993). The improved approach abandons the eight-quarter moving average in favor of weighted regression forecasting, but retains variance specification similar to the current approach, which pools individual item variances, a necessity given the wildness of the VTMR estimator. Figure 2.5 shows a subset of the data in Figure 2.4 in which, for the sake of clarity, we eliminated samples whose VTMRs were 10 or greater.

[10]These procedures are discussed at greater length in Adams et al. (1993).

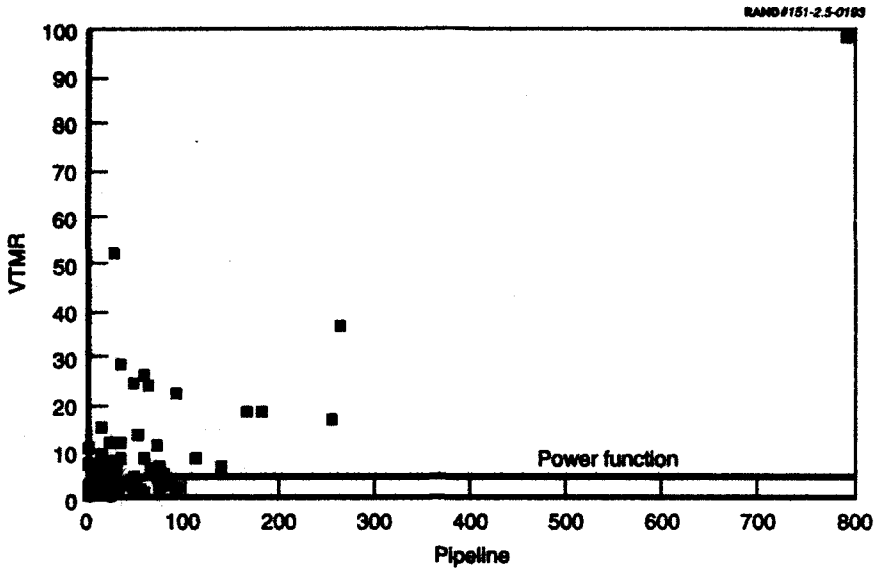

Figure 2.4—Scatter Diagram of VTMRs and Item Pipelines

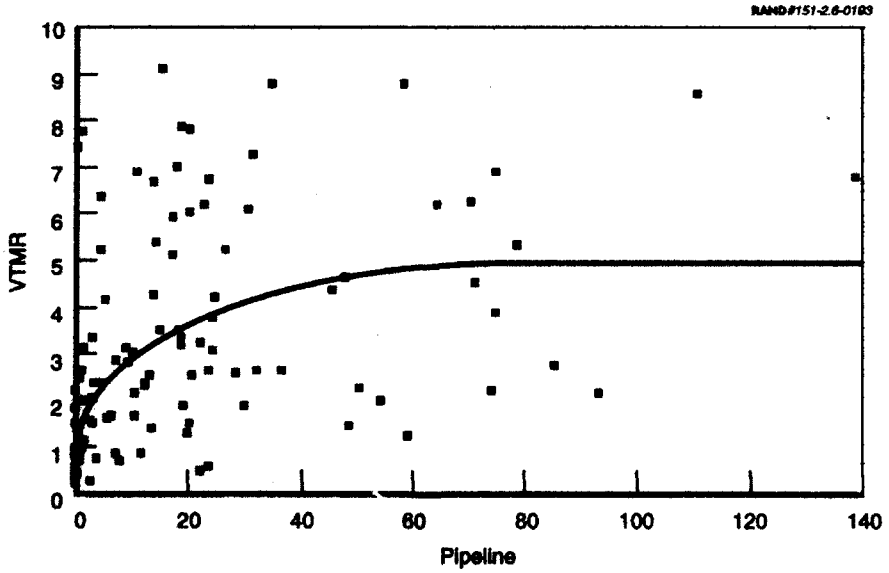

Figure 2.5—Scatter Diagram of VTMRs and Item Pipelines (Enlarged)

Another modeling feature of the current system is that it multiplies together a large number of probabilities to estimate the aircraft availability that a particular mix of stock levels will deliver. For every applicable component, the Aircraft Availability Model estimates the probability that no aircraft is down owing to a shortage of that component and, assuming independence among applicable components, multiplies those probabilities together to estimate the probability that no aircraft is short of any applicable component. Given that an aircraft may have, say, 300 recoverable LRUs, the achievement of an aircraft availability goal of 83 percent requires that virtually every item has sufficient stock that its probability of not being short is virtually one.

As we mentioned before, the current system assumes that the recoverable item inventory system is managed with a continuous review reorder policy and that the quantity per requisition is always one. For some items, the observed reorder behavior is different from the assumed behavior. For these items, requisition quantities exceed one, especially for items that have large quantities per application (QPA). If requisitions for a particular family subgroup arrive according to a Poisson process, then the quantity in the requisition acts as a variance multiplier. For example, if the requisitions were always submitted for quantities of 10, the VTMR of that demand process would be 10, not 1. Thus the underlying Poisson process is masked by the requisition size. This "batching" behavior may lead to concern about high observed VTMRs, but when VTMRs of items that are routinely ordered in batches are judged to be large in some sense, the batching must be accounted for. Generally, large requisition sizes tend to be highly correlated with large QPAs. Thus the demand processes of items such as turbine blades, stator and rotor vanes, guide vanes, and similar items tend to have high VTMRs when their requisitioning behavior may be quite regular.

Assumptions About Management Adaptations

In its computation of safety levels, the current system assumes that parts shortages are randomly distributed among aircraft tail numbers, an assumption that implies little or no cannibalization at the usual levels of system performance. The system also assumes that resupply times are independent of the richness of the asset position. This would imply that, even when faced with a severe paucity of serviceable spares, the system doesn't "hurry," so to speak, an assumption that tends to exacerbate the propensity of the current system to overinvest in spares.

On the other hand, we know that management adaptations are routinely practiced even in peacetime to mitigate the effects of asset shortages. Those adaptations take the form of cannibalization, lateral supply, expedited repair and transportation, and so on. It would be an overstatement to say that none of these adaptations is accounted for in the current requirements estimation process, but it is accurate to say that such adaptations and their associated costs are not accounted for sufficiently, and that they have important effects on system performance. The substantial role of management adaptations in shaping system performance is discussed at length in Section 4.

The Current System's Treatment of WRSK

The current system, as described in Figure 2.3, computes requirements for WRSK assets separately from POS. The POS requirements computation then treats them as additives. POS and WRSK assets are, in a sense, pooled in practice. POS assets could certainly be used in wartime, but no explicit account is taken of them in computing WRSK requirements. For many NSNs, bases have authorized levels of both categories of assets; thus a base's assets of a given NSN may comprise both POS and WRSK, and the assets are accounted for and stored separately. Despite this partitioning, WRSK assets are frequently withdrawn to satisfy urgent requirements when there is no POS asset available at a base to satisfy the need. In this sense, then, WRSK assets and POS assets are pooled in use. This practice is examined further in Section 4.

3. UNCERTAINTY, SYSTEM PERFORMANCE, AND "LONG SUPPLY"

A fundamentally important characteristic of the current requirements system is its view of the world as a steady-state system. This view has three implications:

- We underestimate the effects of uncertainty on system performance.
- We do not forecast outyear requirements very accurately.
- We are more vulnerable than necessary to the future eventuating differently than forecast because we do not explicate uncertainty in our model of the future.

We discuss these implications further in this section, offer an explanation of the underlying causes of the long supply problem as we previously defined it, and suggest some simple strategies for partially mitigating the effects of churn.

UNCERTAINTY

Figure 3.1 illustrates an important characteristic of outyear requirements forecasts based on a steady-state assumption. The reason that outyear requirements are so minuscule is that the current requirements system does not account for churn. In its estimations of outyear budgetary requirements, it only recognizes the need to buy condemnation replacements and respond to any planned increases in program. This phenomenon, observed over the years, may account for what people in the system sometimes refer to as the "bow wave." It results from the assumptions in the current system that (a) item characteristics do not change, (b) the items now in the database will still be in the database in the outyears, and (c) no new items will appear in the database. *While unfunded requirements may from time to time play a role in creating this phenomenon, it typically is not a problem of funding shortfall; it is a problem of reasonableness in assumptions. It may be more important than any other single factor in our inability to forecast outyear spares and repair requirements accurately.*

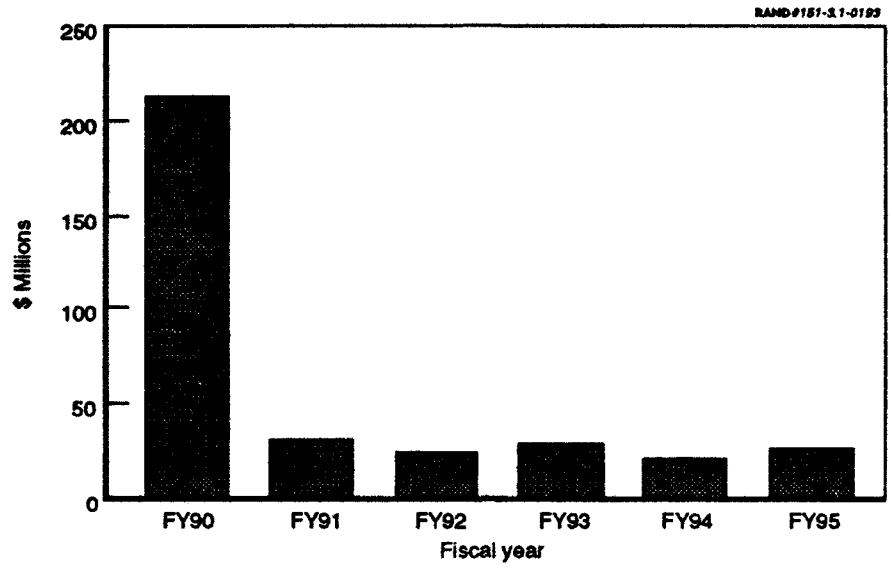

Figure 3.1—A Notional Characterization of Budgetary Requirements Profiles Resulting from the Steady-State Assumption

Obviously, churn has important implications for long supply. In determining whether an item is in long supply, in a buy position, or not, we compare the total number of assets of that type in the inventory system with the number needed a lead time beyond the buy point as estimated by the requirements computation. Call that estimated number N_i for item i. The current asset position of item i was shaped, perhaps, by past values of N_i. For example, if N_i was computed to be 300 four years ago and its current value is only 200, chances are good that the item is in long supply because the item manager would probably have initiated a purchase request four years ago and the assets would probably still be in the system. If numerical values of $\{N_i\}$ are especially volatile, we may be chasing them all the time, but never "get it right" in the sense of asset positions conforming to the current values of the $\{N_i\}$. We did some theoretical explorations of the problem of long supply in which we showed that it is sensitive to demand variability as well as pipeline size.

THE ROLE OF UNCERTAINTY IN LONG SUPPLY

Quite early in the course of this work, we wished to examine the hypothesis that random variation in demands will naturally lead to long supply. We again define long supply as the condition in which there

are more spares in the inventory system than the number the requirements computation estimates as required a lead time beyond the buy point. We tried to formulate analytic models that would enable us to compute the probability of an item being in long supply under various assumptions about the characteristics of the item (its weighted resupply time, demand rate, NRTS rate, etc.). We found the problem intractable, and decided to conduct a simulation experiment.

Our simulation study simplified the problem that one encounters in the real world. We assumed, for example, that there was no safety stock in the system. All that was required was to fill the pipeline. We assumed a known, constant resupply time and zero condemnation rate. We defined long supply very simply as having more spares in the system than the number required to fill the pipeline that was assumed to result from the random generation of quarterly demands. We simulated 80 quarters of demand and replicated each experiment 500 times. In each quarter, we drew at random the number of demands that occurred during that quarter from a probability distribution. We used this number to estimate the current number of spares required to fill the pipeline with no safety stock. The estimated number of spares required to fill the pipeline became, in effect, the N_i, already described, except that, as already mentioned, it did not include safety stock. The simulation study consisted of estimating the probability that an item with certain characteristics would be in long supply across its life in the inventory system. We varied the item's characteristics, using mean quarterly demands of 45.5, 12.13, and 6.07; a constant procurement lead time of 8 quarters; weighted resupply times of 15, 30, and 60 days that correspond roughly to NRTS rates of 25, 50, and 100 percent, and variance-to-mean ratios of quarterly demands of 0.5, 1.0, and 3.0. Initially, stock was set equal to the expected number of assets in resupply (i.e., the expected pipeline size). After each quarter, an estimate was made, using an eight-quarter moving average, of the future pipeline. Additional stock was bought if, and only if, the expected future pipeline exceeded current stock by one or more units. What follows is a presentation of selected simulation results that illustrate the importance of pipeline size and variability in exacerbating the long supply problem.

Figure 3.2 shows the probability that an item is in long supply as a function of demand rate. The probability of being in long supply increases more over time if the item has a high demand rate than if it has a low demand rate. In this example, demand was described as a Poisson process. High, medium, and low demand rates were specified to be 45.5, 12.13, and 6.07 demands per quarter, respectively. Note

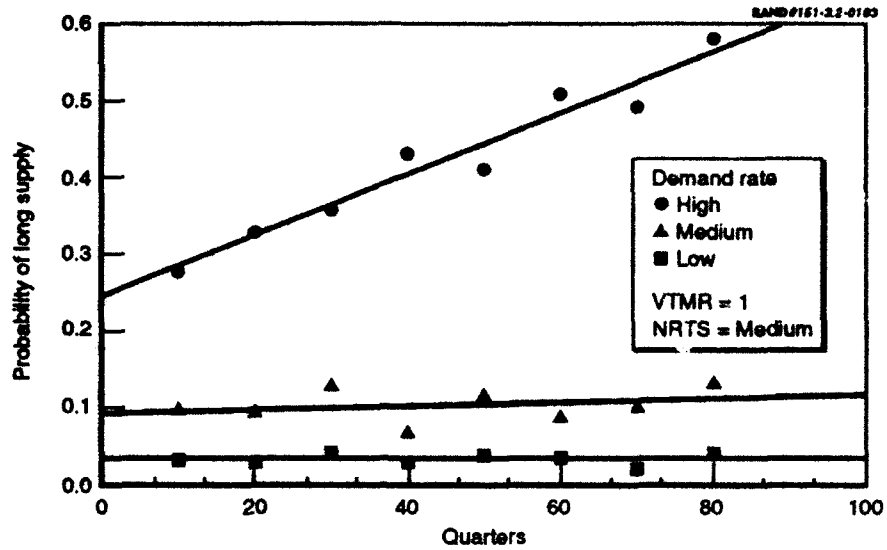

Figure 3.2—Probability of Long Supply vs Demand Rate

that six demands per quarter is not "low" relative to the total spares inventory. It is "low" among higher-demand items, which is the set of interest here. Truly low-demand items, as Figure 3.2 suggests, have a low probability of being in long supply.

Figure 3.3 portrays the probability of being in long supply as a function of the NRTS rate. As the NRTS rate increases, the probability of being in long supply increases more rapidly over time. In this case, we are really observing the effects of a larger pipeline because, given a specified demand rate, a higher NRTS rate results in a larger pipeline owing to depot repair turnaround times being much longer than base repair times. High, medium, and low NRTS rates were defined to be 100, 50, and 25 percent, respectively.

Figure 3.4 shows the probability of being in long supply as a function of demand variability. It is not surprising, perhaps, that the higher the variability, the greater the probability of being in long supply. In this case, the underlying probability distribution is negative binomial.

It is important to note that NRTS and demand rates combine multiplicatively to estimate an item's pipeline. Thus the conclusions about the probability of being in long supply extend directly to the pipeline as well as to NRTS and demand rates alone.

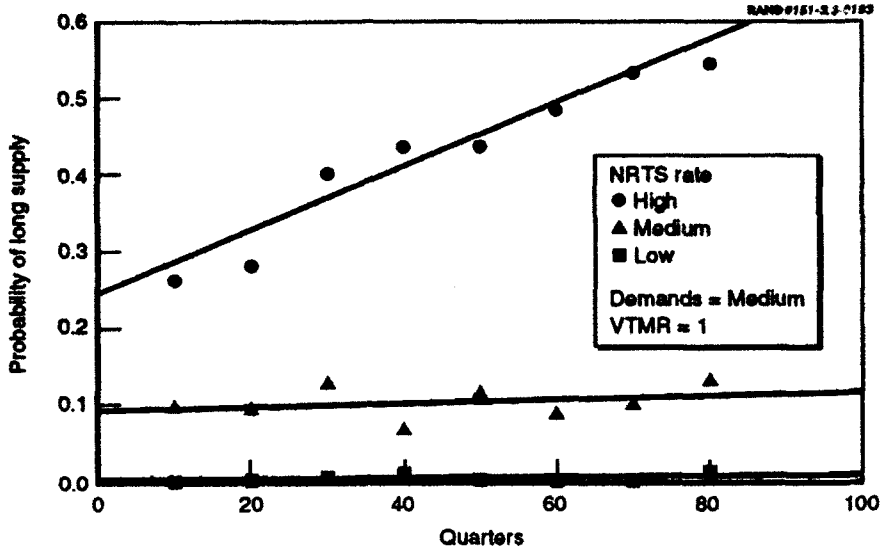

Figure 3.3—Probability of Long Supply vs NRTS Rate

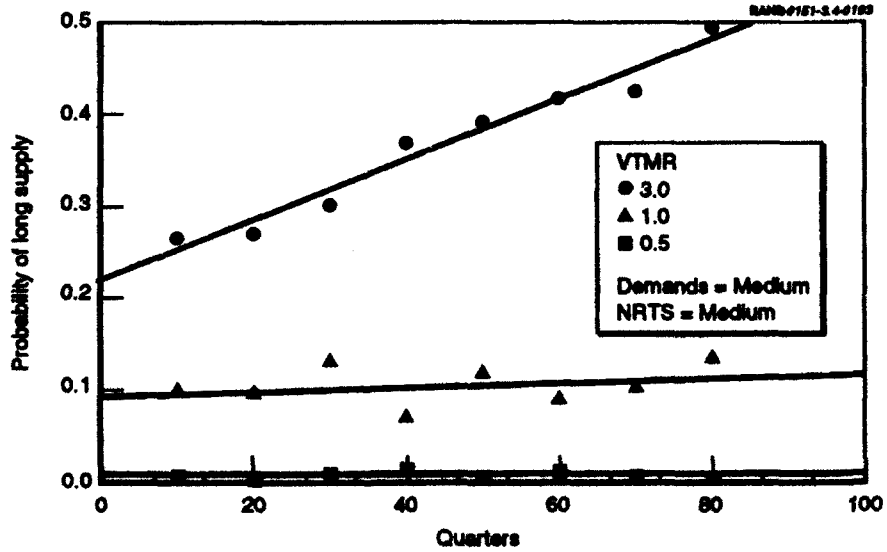

Figure 3.4—Probability of Long Supply vs VTMR

In summary, then, one can conclude that the probability of being in long supply is an increasing function of the pipeline and the variability in the underlying demand process. These findings, we believe, have important implications for strategies to avoid long supply because they identify the item characteristics that tend to be associated with the problem.

Figure 3.5 estimates the probability of being in long supply as a function of increasing and decreasing demand rates. In this case, we specified a linear increase or decrease of 50 percent over the item's 20-year life in the inventory system. These curves pool the results of 500 experimental replications for each case.

With decreasing demand, the item has a steadily growing probability of being in long supply. With increasing demand, the probability varies over time but is never very large. The item portrayed here has a medium demand rate at the start of its life, roughly 12 demands per quarter, and a rather modest pipeline of about four assets; therefore, when a single asset is added to the inventory system, it has obvious effects on the probability of the item being in long supply. These effects are noted here around the 30th and 40th quarters and again around the 70th and 80th quarters. Clearly, in the case of increasing

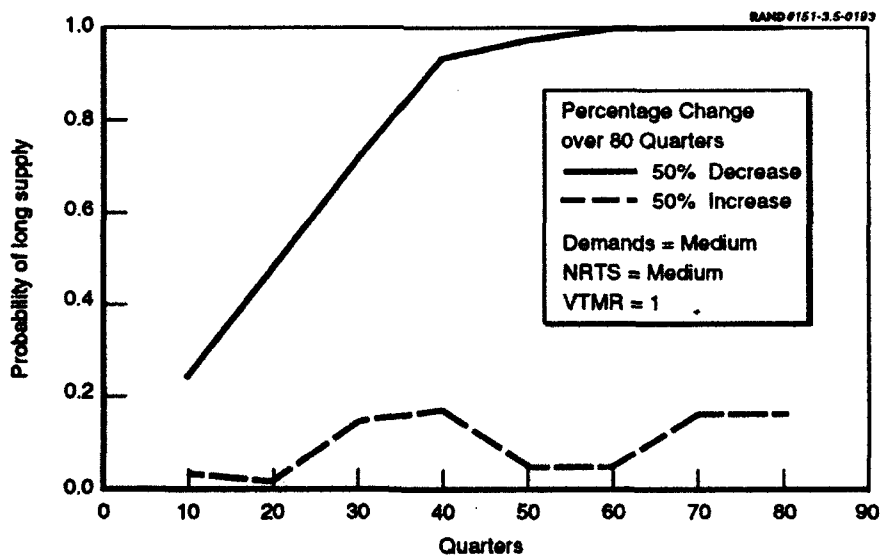

Figure 3.5—Probability of Long Supply vs Changing Mean Demand for a Medium Demand Item

demand, long supply isn't of much concern; however, note that the item demand rate was held linear in these experiments, a rare condition in the actual inventory system. The more important problem with increasing demand is filling the backorders that are likely to eventuate.

The item portrayed in Figure 3.6 has a "high" demand rate, roughly 45 demands per quarter, and a pipeline of about 15 assets. With linearly decreasing demand, the probability of the item being in long supply is quite high even by the tenth quarter, and it grows rapidly. After 30 quarters, it is virtually certain to be in long supply. With increasing demand, the probability of being in long supply remains quite small.

One can summarize these findings quite simply. The probability that an item is in long supply during its active life in the inventory system is substantial; it increases with the size of the item's pipeline and with the variability in its demand over time. Long supply is not a significant problem with small pipelines. The practice of rounding the computed number of assets required upward to the nearest integer may account in part for this observation. Nonstationarity in demand also plays an important role in determining the probability that an item will be in long supply.

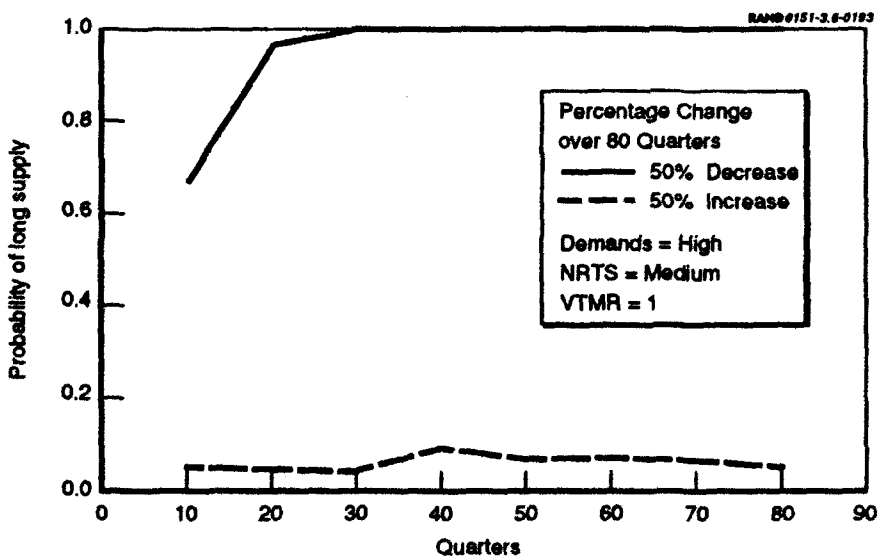

Figure 3.6—Probability of Long Supply vs Changing Mean Demand for a High Demand Item

These figures tell only part of the story. If one examines the requirements database to verify that the items with large pipelines and high demand variability are indeed the items that tend to be in long supply, one may be surprised to find that it is not so. The preponderance of items in long supply have zero pipelines. In examining item demand histories, we observe that the items with the greatest numbers of assets in long supply are the ones that *used to have* large pipelines and high variability. This is one of the effects of churn in the database.

IMPLICATIONS OF THESE EXPLORATIONS FOR MANAGEMENT STRATEGIES

In the next two sections, we explore the implications of management adaptations in spares requirements estimation. Taking explicit account of their effects on the performance of specified stockage postures can help mitigate the long supply problem. An additional strategy that can help to a limited extent in mitigating long supply is that of delaying the specification of spares procurement quantities until the end of the administrative lead time of the procurement process. When initiating a purchase request for spares, there is usually *greater uncertainty* about the quantity needed than there will be six to nine months later, which is the typical length of time involved in the administrative lead time. During the administrative lead time, additional data reflecting demands and NRTS actions will become available. If those data are taken into account, they may have implications for the procurement quantity. Therefore, we suggest that the procurement quantity be reviewed just before the procurement contract is signed to ensure that the most current characteristics of the items being procured are appropriately reflected in the contractual quantity. It is our understanding that AFMC has already implemented this strategy.

Another form of hedging is to understate the procurement quantity deliberately. One could order only enough additional assets to fill the pipeline, or the pipeline plus one standard deviation, say, and negotiate a contingency quantity or option to be specified at some future date. In fact, a combination of these two strategies might be useful in mitigating the effects of additional spares procurements on the long supply problem.

4. UNCERTAINTY AND MANAGEMENT ADAPTATIONS

A principal aim of this research was to explore the effects of uncertainty and management adaptations on spares requirements estimation. Although uncertainty tends to cause the system to deliver worse performance than anticipated in the requirements computation, management adaptations more than overcome it. On balance, the system delivers better performance than the requirements system thinks it is buying because the requirements system ignores many of the management adaptations that are quite routinely practiced, even in peacetime, e.g., cannibalization, priority repair, lateral repair and lateral supply, and expedited handling and transportation.

We devised a system of software that enables us to replicate the current requirements and central stock leveling systems and to evaluate the performance in both peacetime and a nominal wartime scenario of the future stockage posture that is anticipated to result from any specified mix of spares procurements.

STRUCTURE OF REQUIREMENTS EVALUATIONS

We are able to replicate the CSIS computations with software we built at RAND. We also have a copy of the AAM operating at RAND. With these two pieces of software, we are able to compute the POS levels for the system worldwide and, of course, the associated budget level.

The POS levels are then fed into our replica of AFMC's D028 central stock leveling system that we use to allocate the stock levels to all of the bases and the depot. That enables us to establish the anticipated POS asset position. We chose the F-16 weapon system as a case study. It is deployed at more than 50 bases worldwide. The allocation model also accounts for common items, i.e., items that are also applicable to one or more weapon systems other than the F-16.

We compute WRSKs with Dyna-METRIC Version 4, performing essentially the same computation as WSMIS. The sum of these two computations, WRSK and POS, constitute the anticipated asset position or stockage posture that we described previously.

That stockage posture is then fed into Dyna-METRIC Version 6, an advanced, hybrid analytic-simulation model, with which we can eval-

uate the performance of the anticipated stockage posture across a variety of peacetime and wartime operational scenarios. Version 6 is a later version of the Air Force's standard capability assessment model. It is described in a companion report (Isaacson and Boren, 1993).

Analysis of the results of these evaluations enables us to devise alternative approaches to the computation that will yield improved stockage postures whose performances will be more robust in the face of uncertainties. Figure 4.1 is a schematic diagram of the software system used in the evaluations discussed in this section.

Our replication of the spares and repair requirements computation includes the entire aircraft recoverable spares inventory system; however, when we evaluate the performance of the anticipated stockage posture, we do it only for the F-16 weapon system. Because Dyna-METRIC Version 6 evaluations require substantial amounts of computer time, we reduced the size of its database by including only those items that affect the outcomes. Toward this end, we made a peacetime and wartime run and selected a set of items that most affected system performance, then reran both cases with only the selected

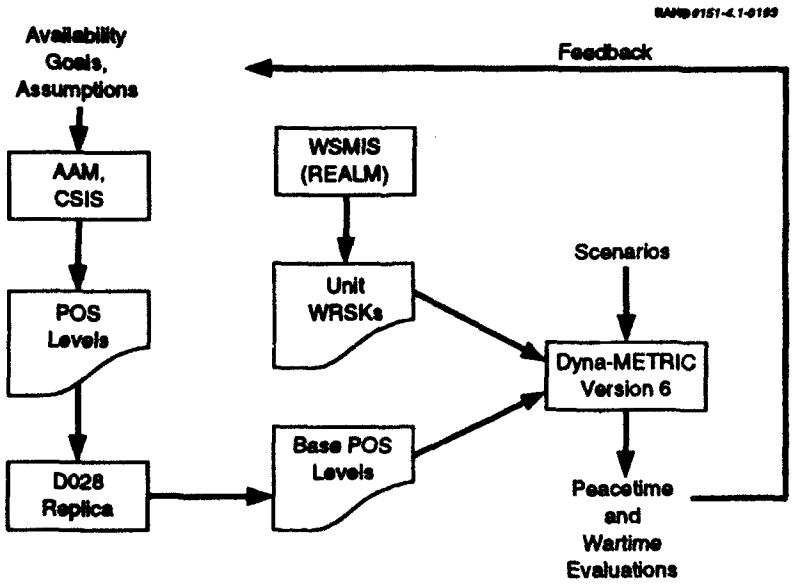

Figure 4.1—Schematic Representation of Evaluation Software System

subset and made sure that the resulting aircraft availability remained unchanged. This process resulted in the inclusion of 646 items in the Dyna-METRIC database, of which 467 were LRUs.

THE QUICK RESPONSE OPTION

In the figures that follow, you will note a management adaptation that we call *quick response*. This term is explained in Table 4.1. The quick response option represents prioritized handling of assets in all phases of the depot-repair pipeline with the exception of the number of days actually spent in repair, i.e., the shop flow days.

Note that the quick response option is assumed to diminish the total depot repair turnaround time by roughly 45 percent. As we will show, it has a powerful effect on system performance.

Figure 4.2 shows our estimates of the peacetime performance of the current requirements system's anticipated stockage posture using the March 1986 D041 database. The evaluations were done using the same estimates of item characteristics (demand rates, NRTS rates, repair times, etc.) reflected in the database used to compute the requirements. Thus no account was taken of the fact that when AFMC computes the requirement, it has no way of knowing how item characteristics will really eventuate three years or so in the future when the assets are received. Thus, these are churn-free estimates.

Table 4.1

The Quick Response Option

	Pipeline Lengths	
	Actual	Quick
Base processing days	4[a]	3
Retrograde days	16	5 overseas 2 CONUS
Supply to maintenance	3[a]	1
Shop flow days	40[a]	40
Serviceable turn-in days	5[a]	2
Order-and-ship time	21	5 overseas 2 CONUS[a]
Total	89[a]	56[a] overseas 48[a] CONUS

[a]Varies by item; average shown.

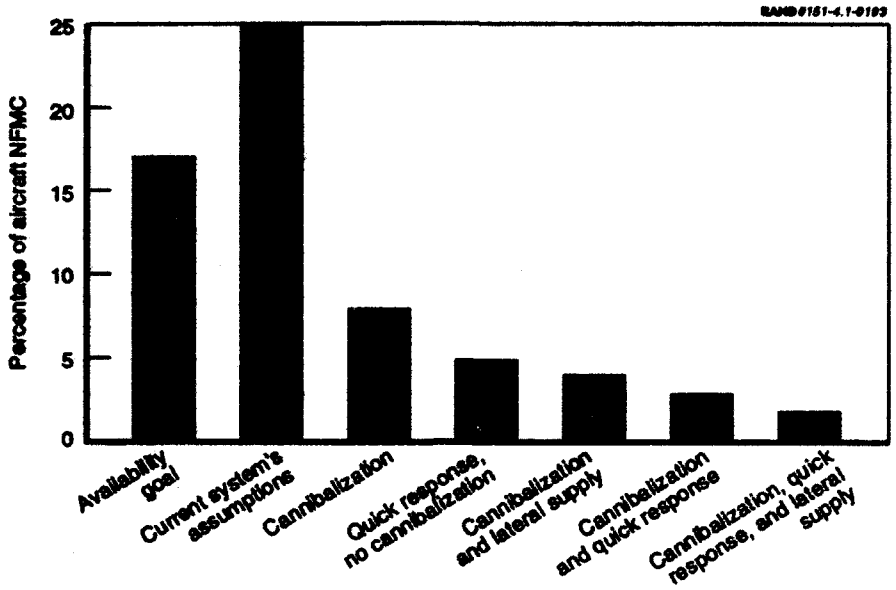

Figure 4.2—Peacetime Performance with WRSK Withdrawals

The leftmost bar in Figure 4.2 reflects the aircraft availability goal of 83 percent that was specified to the requirements computation, i.e., 17 percent of the force not fully mission capable owing to aircraft waiting for parts to be repaired or shipped to the location where needed. The bar labeled "Current System's Assumptions" is Dyna-METRIC's estimate of the peacetime performance of the anticipated stockage posture produced by the requirements computation using the March 1986 database. It is worse than the specified goal because of an assumption made in Dyna-METRIC that items whose application percentages are less than 100 are used in equal proportion on the aircraft at every location, thus negating economies of scale in safety stock. Unfortunately, there is no way to avoid this assumption given the data available. Dyna-METRIC Version 6 also assumes cannibalization of all items, and consolidates shortages perfectly, thus inducing an estimating error in the opposite direction. We do not claim that these errors are offsetting.

Note how dramatically management adaptations of the variety evaluated here improve system performance. Cannibalization or the quick response option alone achieves better than 90 percent availability. Lateral supply improves performance even more. The three adapta-

tions together achieve 98 percent aircraft availability, far more than the specified goal. Recall, however, that these evaluations essentially assume perfect information in the sense that the same item characteristics are assumed to apply to the point in time in 1989 when this asset position is assumed to eventuate as were assumed in the requirements computation itself.

In Figure 4.3, the dark gray bars are the same as those in Figure 4.2: They represent the estimated performance of the anticipated asset position using the same item characteristics in the evaluations that were used in the requirements computation. The light gray bars in this graph portray the estimated performance using item demand rates that were observed at the time that the asset position would have evolved in the system, roughly three years after the computation. The variance of the item pipelines was specified to be the same as in the requirements computation. The estimated system performance would have been worse had we used the demand variances actually observed in 1989. Performance is significantly degraded owing only to our inability to forecast the future. This is what churn in the database really means: We will not achieve performance as good as

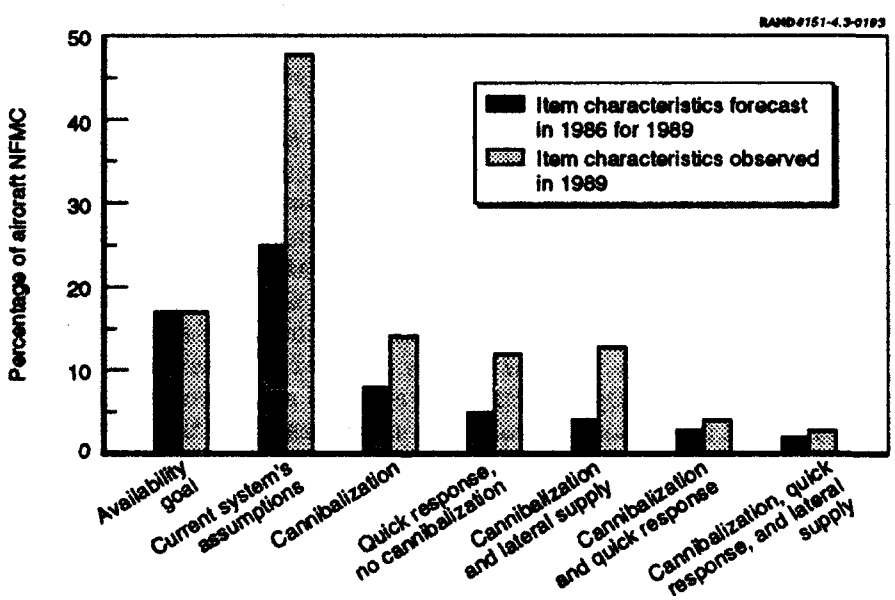

Figure 4.3—Effects of Churn on System Performance with WRSK Withdrawals

we anticipate because we fail to account for the state-of-the-world and statistical uncertainties to which the system's performance is vulnerable.

If we made the same naive assumptions about the operation of the logistics system as are made in the current requirements computational system, we could expect to see the performance portrayed by the bar marked "Current system's assumptions." On the other hand, the logistics system doesn't really perform the way the requirements system assumes that it does. In reality, management adaptations are in place and to some extent are routinely practiced even in peacetime. With such management adaptations systematically applied, the performance of the logistics system, even in the face of realistic levels of churn in the database, is dramatically improved, almost to the levels achieved assuming no churn at all.

Our evaluations overstate the benefits of cannibalization, since DynaMETRIC Version 6 assumes perfect consolidation of shortages of all items at a base, an assumption inconsistent with reality. Real-world performance is probably close to the case labeled "Cannibalization and lateral supply." F-16 not mission capable supply (NMCS) rates at the time of these observations were around 9 percent. In the face of such performance, availability goals of around 83 percent were persistently specified to the requirements computation. It is simply a numeric artifact of the aircraft availability model that such goal specification can be expected to produce much higher levels of availability in real life because management adaptations do have such powerful effects on system performance, especially in overcoming uncertainty.

In Figure 4.4, we add another set of bars, the white bars, to examine the effects of an alternative WRSK withdrawal policy, i.e., the WRSK is held inviolate during peacetime; no WRSK withdrawals are allowed. Note that this policy has little effect on system performance. Examination of the distribution of MICAP[1] termination codes suggests that WRSK withdrawals are frequently used during peacetime to satisfy MICAP requirements for serviceable spares. A WRSK withdrawal, even with the approvals required, is fairly quick and convenient. It obviates the need for the extra work implicit in cannibalization and its associated risk of damage and is quicker than a lateral supply action. If WRSK withdrawals are prohibited, however,

[1]Shortages judged to affect the mission capability of an aircraft.

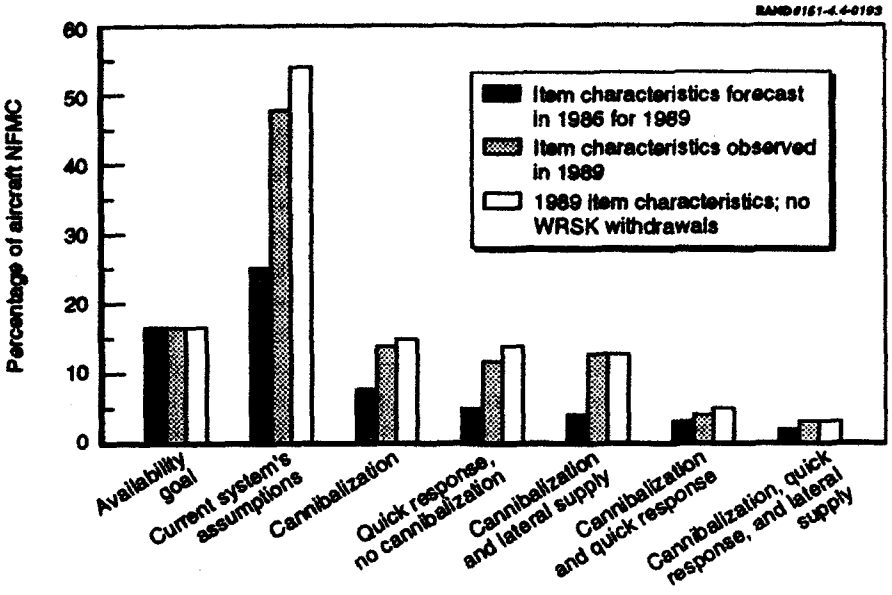

Figure 4.4—Effects of Churn on Peacetime Performance, No WRSK Withdrawals Allowed

lateral supply and cannibalizations would then be brought into play more frequently, thus avoiding serious degradation in performance, although costs would clearly rise.

Figure 4.5 reflects some assessments of wartime performance. The measure of effectiveness is the proportion of aircraft not fully mission capable at the end of 30 days of war. The war is assumed to occur after 300 days of peacetime operation to ensure minimal effects of initial conditions. Full cannibalization is assumed in each of these cases. In these evaluations, the same assumptions were made as in Figure 4.2: We assume perfect knowledge of item characteristics in the future when we do the requirements computation. The current performance is roughly the same (within experimental error) as anticipated in the computation of the WRSK. Quick response and lateral supply again improve performance dramatically. Even when there is no depot replenishment in the first 30 days of wartime, the quick response option pays off because the units go to war with fuller WRSKs, i.e., there are fewer assets tied up in depot repair pipelines, and those assets are now serviceable.

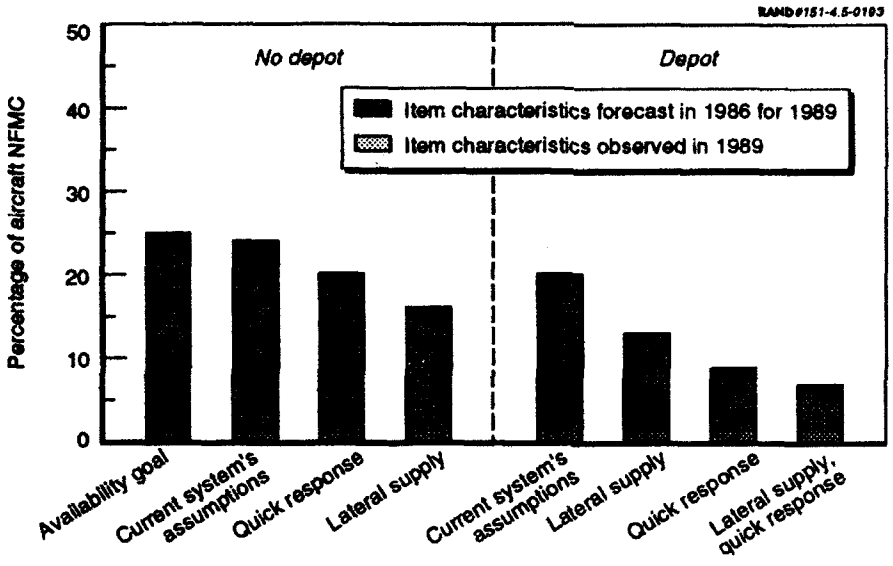

Figure 4.5—Estimated Wartime Performance, WRSK + Primary Operating Stock

In the case where depot replenishment is assumed to take place in the first 30 days of war, the quick response option has much greater payoff. When coupled with lateral supply, it achieves a very impressive level of performance when we do not account for churn in the database.

Taking account of churn, however, changes the picture considerably, especially without depot replenishment. As Figure 4.6 shows, depot replenishment, coupled with the other management adaptations, helps the system overcome the effects of churn.

The evaluations shown in Figure 4.7 explicitly examine the effects of prohibiting WRSK withdrawals on both peacetime and wartime performance. Without cannibalization, denying access to the WRSK in peacetime degrades system performance. The percentage of aircraft NFMC increases from 25 to 31 percent. Note, however, that cannibalization alone improves the performance significantly. With cannibalization in place, the policy of prohibiting WRSK withdrawals only increases the percentage NFMC from 8 to 10.

The bar on the far right of Figure 4.7 reflects the wartime performance when WRSK withdrawals are allowed during peacetime, but

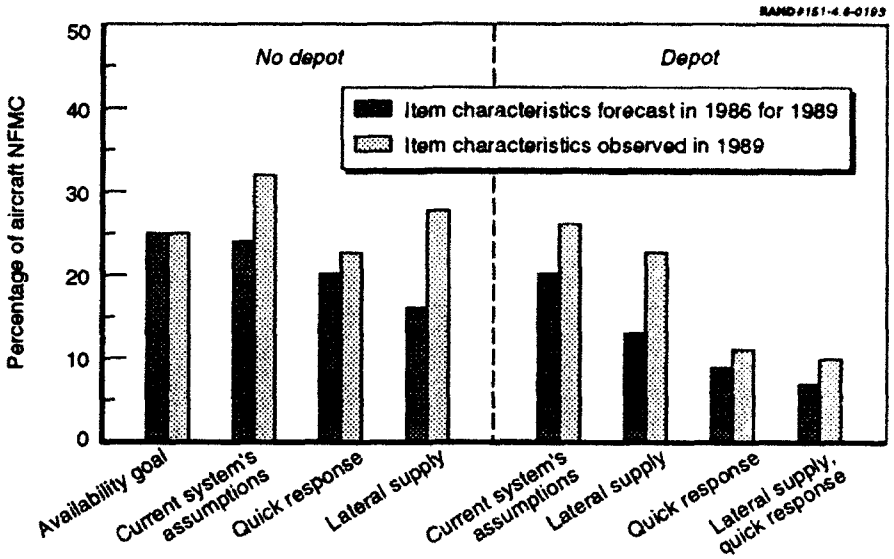

Figure 4.6—Effects of Churn on Wartime Performance, WRSK + Primary Operating Stock

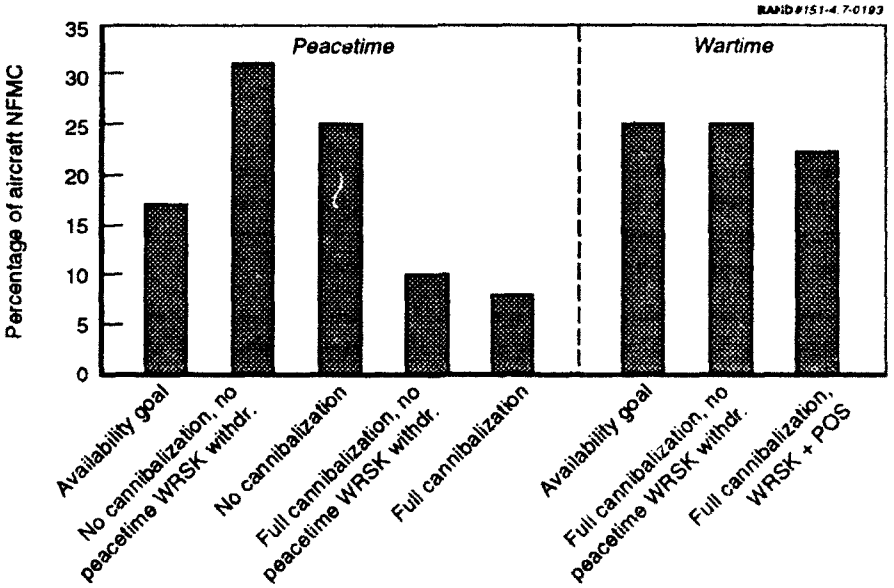

Figure 4.7—Performance of Alternative WRSK Withdrawal Policies

POS assets are deployed along with the assets left in the WRSK on D-day, tending to offset the degradation in wartime performance that might otherwise have resulted from WRSK withdrawals in peacetime. The evaluations portrayed in Figure 4.7 do not account for churn.

Figure 4.8 reflects the effects of churn. In Figure 4.8, the dark gray bars are the same as those in Figure 4.7. The light gray bars in Figure 4.8 reflect the effects of churn. As before, management adaptations are very effective in mitigating its effects. The results in Figure 4.8 suggest that the effects of deploying certain POS assets in wartime can more than offset the effects of WRSK withdrawals in peacetime. Clearly, constraints in transportation capacity would need to be accounted for to include some POS assets in the airlift-deployable WRSKs.

IMPLICATIONS OF THE EVALUATIONS FOR SPARES REQUIREMENTS

In summary, the evaluations show that management adaptations of the several varieties examined here more than overcome the degrading effects of the state-of-the-world and statistical uncertainties that

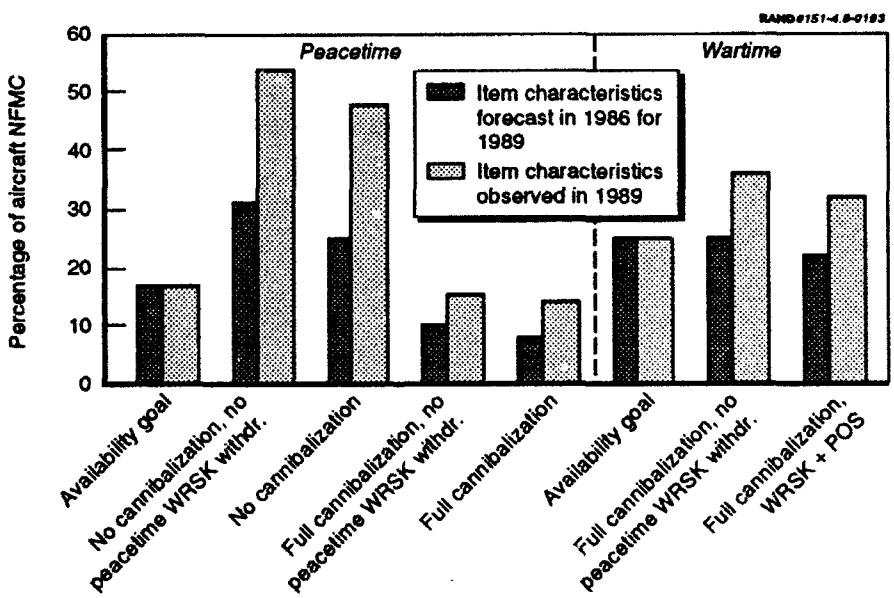

Figure 4.8—Effects of Churn on Performance of Alternative WRSK Withdrawal Policies

make the spares requirements estimation problem so difficult. Many of the features of the current requirements system tend to induce a richness in spares investments. For example, the computational model, with traditional aircraft availability goals, buys safety stock to achieve levels of item availability that are virtually one.[2] Several other features of the current system result in the propensity for overinvestment. It seems clear that a leaner mix of stock levels in the system could easily be accommodated while maintaining desired levels of performance. In Section 5, we discuss some of the initiatives that we believe would contribute to that goal and demonstrate their effectiveness.

Conspicuous by its absence from the discussions in this section is the prioritization of repair at both base and depot level. Abell et al. (1992) showed that priority repair has very dramatic effects on aircraft availability when applied to depot-level component repair. Base-level priority repair also has important effects, and neither of those effects has been accounted for in these evaluations. Thus, in this sense at least, the evaluations are somewhat conservative, although the full cannibalization assumption of which we have included the effects is, as a practical matter, not achievable in the real world. These two assumptions, full cannibalization and the omission of priority repair, act in opposite directions, but we do not have the data to quantify their effects.

[2]Item availability is defined as the probability that an aircraft selected at random will *not* have a shortage of the item.

5. TOWARD IMPROVED ESTIMATION OF SPARES REQUIREMENTS

The topics and ideas we discuss in this section emerge from three sources: (a) the features of the current system and its logical inconsistencies with the actual logistics system, (b) the implications of the evaluations presented in Section 4, and (c) other conclusions reached in this work and in related work by others. Some of these topics and ideas are developed and discussed at length in companion reports. The following list summarizes them, but they are not sequenced in any particular way. They may seem somewhat unrelated, but they share the common characteristic that they all affect either the level of investment in aircraft recoverable spares or the cost-effectiveness of the mix of spares procured. We discuss each in turn:

- Take aggressive steps to enhance the responsiveness of depot-level component repair to reduce pipelines of selected items.
- Implement a computational model that provides for cannibalization of designated items.
- Implement improved demand modeling and forecasting techniques coupled with a Poisson assumption to determine investment levels.
- Add engines and engine modules to the D041 database.
- Conduct an intensive, continuing, in-depth data and file maintenance training program for IMs and equipment specialists and correct the data processing logic used to produce the indenture file.
- Develop an aircraft configuration database.
- Given the required configuration database, evaluate the concept of consolidating WRSK centrally and computing, assembling, and deploying WRSKs as needed to support deploying units.

We discuss each of these ideas in the pages that follow and, where appropriate, refer the reader to other documentation for more complete exposition.

ENHANCING THE RESPONSIVENESS OF DEPOT REPAIR

One implication of the evaluations in Section 4 is clear: *reducing the depot turnaround times of selected items has powerful effects on system performance and, therefore, on the required levels of spares invest-*

ments given specified aircraft availability goals. In every case we examined, the payoff of the quick response option was dramatic, but its costs also need to be assessed. While enhancing responsiveness would not improve our estimation of spares and repair requirements, it certainly would have an important effect on the level of investment required to achieve specified levels of aircraft availability; therefore, we include our discussion of it here.

In other ongoing work at RAND, additional evaluations and demonstrations of initiatives that are intended to make the depot component repair system more responsive to the most urgent needs of the combat force are being examined. Those initiatives include:

- Development and demonstration of an integrated, multi-echelon spares and repair management system,
- Demonstration and evaluation of proactive SRU repair, a version of a just-in-time inventory system for recoverable SRUs in support of LRU repair, and exploration of improved consumable stockage policies for depot repair,
- Extension of the logic underlying DRIVE to the prioritization of transportation to achieve aircraft availability goals, and
- Development and demonstration of a decision support system for depot repair planning that will enhance our ability to estimate the "best" mix of component repairs over a planning horizon to achieve specified aircraft availability goals at least cost, taking explicit account of budgetary and capacity constraints, estimated aircraft availability that will be achieved, and the resulting operating position of the stock fund.

The feasibility of several of these initiatives has already been assessed. Given the dramatic effects of enhancing depot repair responsiveness we have seen in the evaluations shown in Section 4, AFMC should be pursuing these initiatives aggressively.

DESIGNATED CANNIBALIZATION

The Air Force has already implemented a coding scheme for WRSK items specifying whether the item is difficult to cannibalize or not. This idea could be extended to items not included in the WRSK, since cannibalization is routinely practiced in peacetime. As we saw in Section 4, cannibalization is a powerful management adaptation for coping with shortages. Ignoring it in spares requirements estimation causes the system to treat items with large quantities per application

the same as items with one per application, and items that are easily cannibalized the same as items that may be very difficult to cannibalize, i.e., it causes resources to be allocated to the wrong mix of spares, given that cannibalization occurs quite routinely during peacetime.

As a matter of policy, the Air Force has traditionally prohibited an assumption of cannibalization for the purpose of computing spares requirements. In an austere funding environment, however, the argument that cannibalization is too costly is considerably weakened; moreover, the ability to apply it selectively only to designated items weakens it further. Given seriously limited resources, one needs to find the most cost-effective mix of spares in which to invest. Designated cannibalization, coupled with a computational model that accommodates it, seems like a reasonable strategy for enhancing the cost-effectiveness of the spares mix, especially in a fiscal environment that is seriously constrained.

Figure 5.1 shows our estimates of the effects of implementing designated cannibalization for an arbitrarily specified case. The computational model used for these computations is described in detail in Gaver et al. (1993). We arbitrarily selected certain federal stock classes as "cannibalizable," so all LRUs and SRUs in those, and only those, federal stock classes were treated as cannibalizable in the evaluations. The components designated as cannibalizable represented about 58 percent of the total cost. The evaluations here are similar to those described in Section 4 except that we used actual item variance-to-mean ratios (VTMRs) observed in 1989 to do the evaluations. We capped these VTMRs at 15.0 to avoid data errors; about 5 percent of the VTMRs were affected by the capping. We followed this procedure in all of the evaluations presented in the remainder of the report. We used the March 1986 database again to compute the spares mix and evaluated the anticipated stockage posture with Dyna-METRIC Version 4 (at the time of this work Version 6 allowed only full or no cannibalization). Version 4, however, does not allow lateral supply; therefore, the only management adaptation we could add here was the quick response option described earlier. Figure 5.1 shows the results. Aircraft availability goals were held constant across the three cases. The results are based on item characteristics as they eventuated in 1989. The costs shown for each policy are for safety stock only.

The assumptions made here could only approximate the effects of a policy of estimating spares requirements using a policy of designated cannibalization. The designation of cannibalizable items was quite

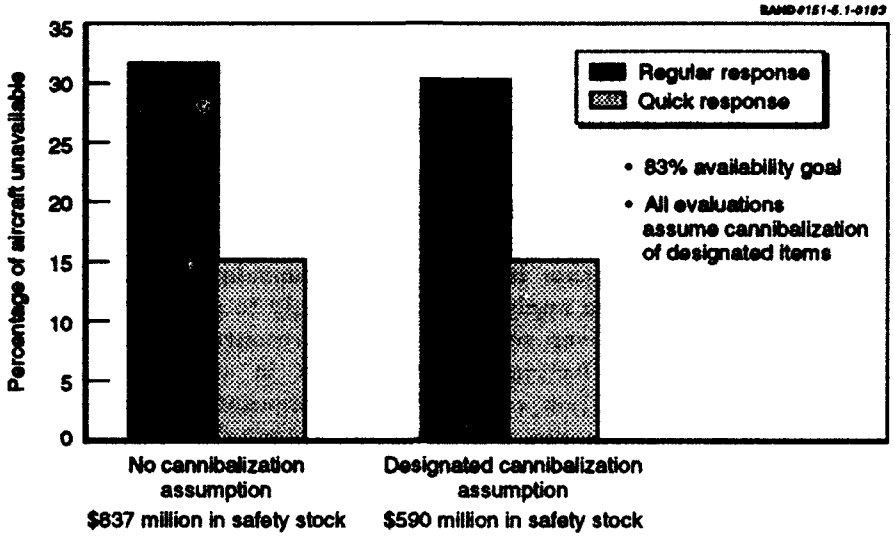

Figure 5.1—Designated Cannibalization Pays Off

arbitrary simply because we had no database to support such an evaluation. We also do not know whether 58 percent of the total cost for designated items is near the actual percentage that would result from item designations by persons competent to make such judgments. Having said that, however, we hasten to add that a policy of designated cannibalization seems to reduce cost substantially without degrading system performance. In our judgment, it is a policy worth pursuing. It seems to have fairly low risk because one can judge the designation of an item as cannibalizable on the basis of its cost in labor hours, its difficulty, and the probability of cannibalizing it successfully without damage. It is simplistic to say that fuses, for example, should be designated items and wing spars should not. The problem is that essentially the entire inventory falls between the two. Although not an idea evaluated in this research, it may be feasible to assign a number between 0 and 1 to designate the "degree of cannibalizability." The simple scheme employed here, as we said earlier, is strictly a 0-or-1 approach.

IMPROVING DEMAND MODELING AND FORECASTING

The current system uses an eight-quarter moving average to estimate item demand rates and NRTS rates. The variance of the probability distribution describing the number of items of each type in resupply is estimated using the observed VTMR of past demands. Demands

are assumed to be strictly proportional to flying hours (for flying-hour-driven items), although we have not seen data that confirm this relationship. The assumption is made that demand processes are stationary, compound Poisson processes. They are, in general, not compound processes but, rather, nonstationary processes (Adams et al., 1993). The current procedure of specifying the variance ignores forecasting uncertainty, accounting only for the stochastic variability in demand. It is approximately correct numerically, but quite by chance.

The eight-quarter moving average is not sufficiently sensitive to nonstationarity in item demands. When an anomaly occurs, it takes two years before its effects vanish. Moreover, it gives as much weight to data several quarters old as to very recent data.

We explored several alternative approaches to forecasting demands and specifying the variance of the pipeline distribution. The improved methods we chose, weighted regression forecasting and variance specification that explicitly estimates forecasting uncertainty as a function of planning horizon length, reduce the demand forecasting error by roughly 40 to 50 percent on items with 15 or more demands per quarter over planning horizons of interest in the requirements problem. Table 5.1 shows the improvement in both root mean squared error (RMSE) and mean absolute deviation (MAD). The improved methods are described in detail in Adams et al. (1993).

The use of the improved methods of forecasting and variance specification resulted in an estimated cost reduction of almost a quarter of a billion dollars ($239 million) in the requirements computation done with the March 1986 database while achieving a level of performance almost equal to that of the current system. Table 5.2 reflects the results of our evaluations of these methods with Dyna-METRIC Version 6. War readiness spares were not included in these evaluations.

It should be noted that forecasting errors degrade system performance in much the same way as data errors and churn do. They

Table 5.1

Percentage Improvement in RMSE and MAD of Improved Techniques over Current System

Measure	10-Quarter Horizon	13-Quarter Horizon
RMSE	48	38
MAD	51	45

Table 5.2

Performance of Improved Techniques Compared to Current System

Management Adaptations	Percentage of Aircraft Unavailable, Peacetime: Current System, $3709 Million	Percentage of Aircraft Unavailable, Peacetime: Improved Methods, $3470 Million
No cannibalization	74.9	76.3
Full cannibalization	33.0	33.1
Cannibalization, lateral supply	17.3	17.2
Cannibalization, quick, lateral supply	3.2	3.6

cause the system to anticipate conditions that do not eventuate. Such errors cost money, just as churn and data errors do, and they diminish the cost-effectiveness of our spares investments. Improvements in forecasting, like correction of data errors, help diminish the levels of uncertainty surrounding our resource allocation decisions.

ADDING ENGINES TO THE REQUIREMENTS DATABASE

In our initial formulation of a base case using the March 1989 D041 database, we noted that engine components represented a very substantial 21.4 percent of the dollar value of F-16 aircraft recoverable organizational and intermediate maintenance (OIM) pipelines. The D041 database does not contain aircraft engines simply because they are not stock-listed items; therefore, engine components are treated by the AAM as aircraft LRUs, since there is no component in the database at the LRU level of indenture for these components and the logic of the program that creates an indenture file from the application file promotes SRUs to LRU status in the absence of the LRU record. These promotions induce overinvestment because shortages of these components are then incorrectly assumed to result in holes in aircraft rather than holes in LRUs.

In its report of a study done several years ago on the effects of including engines in the requirements database, the Logistics Management Institute (LMI) provided a persuasive argument in favor of the idea (King et al., 1986). LMI cites "a 20 percent reduction in the safety level requirement for engine components of" the C-141, F-15, and F-16 aircraft used in its case study of the issue. The idea here is not to compute requirements for engines in D041, only to include a representation of engines in the D041 database to support a more accurate computation of requirements for engine components. LMI's logic is persuasive, and the magnitude of the effects could be significant; certainly engine component pipelines are a significant proportion of to-

tal pipelines. We have nothing else to add to LMI's conclusions, except to register our agreement with them. Recently, engine modules were added to the stock list and AFMC has added them to the database. Engines should be added as well.

DATA AND DATA PROCESSING ISSUES

The Air Force is well aware of the problem of inaccurate data that seems to plague the process of estimating requirements. A major study undertaken several years ago, popularly known as the *Dirty Data Study,* examined several aspects of data accuracy in reporting, processing, and transmission, and found that errors were frequent and significant in their effects, especially errors of omission that induced underestimation of component demands (Coffman et al., 1988). Later follow-up found significant improvement but not complete resolution of the problem (Greenlaw et al., 1989).

In our research, we discovered data and data-processing problems, especially in the creation of the indenture file, that we felt were significant enough to warrant more extensive treatment in a separate report (Abell and Finnegan, forthcoming). As we point out in Section 6, the problem of data quality also has important implications for estimating component repair requirements. We conclude that a massive, intensive, continuing training effort is needed to correct the problems in the D041 database and maintain the accuracy needed to support spares and repair requirements computations. Such a training program should be accompanied by audits of sample data files pertaining to specific LRU families, i.e., LRUs and their indentured SRUs, to (a) understand the current level of data quality through this sampling scheme, and (b) identify specific item managers and equipment specialists who need assistance or training.

DEVELOPING A CONFIGURATION DATABASE

In Section 3, we discussed the problems associated with the number of users in the current system. This number is used to partition item pipelines. It is simply the number of bases with two or more demands in the past 12 months. It is used in lieu of data that would accurately identify an item's programmed exposure to demand at locations worldwide through the use of aircraft configuration, force beddown, and associated program data. The aircraft configuration database needed would contain a mapping of subgroup master stock numbers to aircraft tail numbers. The Air Force has purchased such configuration databases for its more recent weapon systems but has not always

provided for their maintenance. They could be routinely updated by transactions that reflect technical order compliances (TOCs) and implementations of engineering change proposals (ECPs) on specific tail numbers. All that would be needed is a mechanism for overlaying ECP and TOC data into the database. The data would need to include revisions to the mapping of subgroup master stock numbers to tail numbers implicit in the TOC or ECP. It is a straightforward data-processing problem.

We are unable to assess the effects of moving to a system that incorporates the use of such data simply because we have no way to estimate the effects of having correct base-specific application percentages for items whose true percentage is less than 100. In an attempt to understand this issue better, we made four evaluation runs, all with the same aircraft availability goals. In the first, we ran the current system and evaluated the anticipated stockage posture. In the second, we modified the requirements computation to use the number of bases as the number of users as we do in the evaluation model. The second run yielded a higher investment level because the specified availability goals were the same, but demands were being projected to occur at more bases. Next, we reran the current system using the investment level that resulted from using the number of bases as the number of users. Finally, we reran the modified computation using the lower investment level that resulted from the current system. (In other words, we reran the second case using the budget level from the first case, and reran the first case using the budget level from the second case.)

The results are shown in Figure 5.2. They are persuasive. For either budget level, getting the number of users "correct" in the sense that it is consistent between the requirements computation and the evaluation model (our surrogate for the real world) makes a substantial difference in aircraft availability of 9 percent of the possessed aircraft. Certainly a configuration database is needed in execution as well as in requirements estimation, for example, in prioritizing the allocation of assets among locations worldwide to achieve specified aircraft availability goals. In requirements estimation it complicates the problem; however, the number of users as currently defined understates the true number and probably also understates the effects of lateral supply. In our judgment, therefore, pursuit of a sound aircraft configuration database, coupled with explicit accounting for lateral supply, would be very worthwhile. It is clear that a more cost-effective mix of spares results from taking explicit account of all locations at which items are exposed to programmed OIM demand.

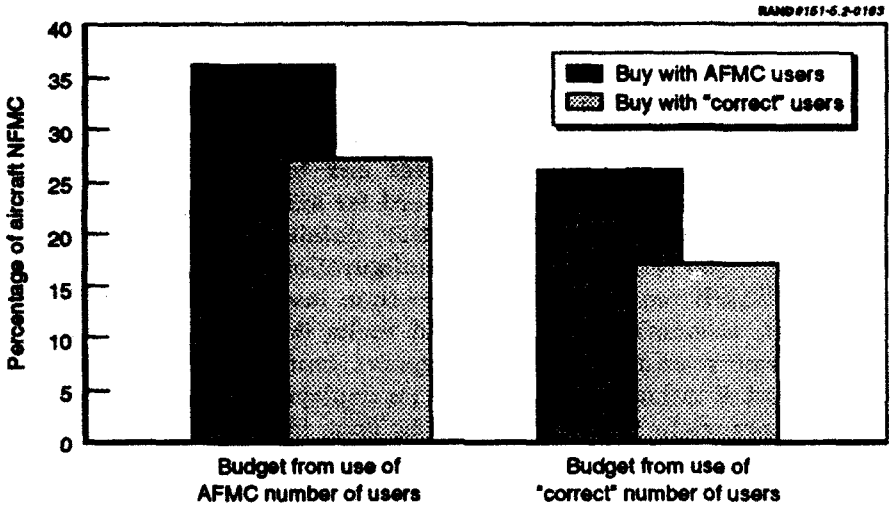

Figure 5.2—Effects of Consistency in the Number of Users

As we point out in the following discussion, such a database would also support consolidation of WRSK at a central location and yield *substantial cost reductions* in WRSK investment. Moreover, the worth of an aircraft configuration database extends beyond its usefulness in supporting the consolidation of war readiness spares. It would support the computation of base-specific application percentages for DRIVE as well as the requirements system and also support the central allocation of stock levels to achieve specified aircraft availability goals without relying solely on demand-based criteria as D028 does. The use of the actual number of bases at which an item is exposed to OIM demand (in lieu of the number of users in D041) and the planned force beddown specified by aircraft tail number and base-specific application percentages (in lieu of the average base assumption), coupled with an explicit model of lateral supply, would increase the cost-effectiveness of the anticipated stockage posture.

CONSOLIDATING WRSK

An alternative policy for providing war readiness spares to deploying units is to maintain them centrally and develop the ability to determine and deploy the appropriate mix of spares with the deployment of any specified mix of aircraft. Current policy allocates WRSKs to every unit with a wartime deployment tasking. While this may have been consistent with a NATO scenario, it seems inconsistent with the

changing global threat and the increased focus of attention on regional deployment contingencies.

The mix of spares in the WRSK is computed by WSMIS/REALM[1] and is adjusted during WRSK reviews conducted jointly by AFMC and the major commands to tailor the spares mix to the specific aircraft assigned to particular units. The need for such a process would be obviated by the aircraft configuration database we have already described. The result of having a configuration database would be a reduced WRSK requirement. Rather than investing in a WRSK for each unit, economies of scale would make it possible to invest in enough war readiness spares to support contingency deployments. Given a set of tail numbers specified for deployment, the computation of the appropriate spares mix could be done in a matter of minutes. The spares would then be pulled from storage and shipped to the deployed unit. The time required to assemble the required WRSK could be minimized by maintaining a generic kit by aircraft model-design-series (MDS) and replacing only those items that did not apply to the particular tail numbers deployed. An exception listing could easily be produced by the computational system.

The F-16 again provides an interesting example. The current WRSK authorization for F-16 units worldwide, for all series, includes 50 individual WRSKs intended to support deployments of from 6 to 48 aircraft, plus 65 packages of war readiness spares to support the LANTIRN, ALQ-119, ALQ-131, ALQ-184, and QRC-80.01 systems. The total worth of these spares is about $604.1 million (based on the costs reflected in the current WRSK database maintained by Headquarters, AFMC). If one were to invest in enough spares to support a deployment twice the size of Operation Desert Shield/Storm (ODS) with 30 days worth of spares usage at 100 percent NRTS rates, the total cost would be about $276.8 million, a saving of $327.3 million, just for the F-16 weapon system.

With such a database, a system of consolidated WRSK storage, maintenance, computation, packaging, and deployment could become a reality and would yield significant personnel reductions at base level as well as substantial reductions in WRSK investment. System performance without war readiness spares at bases would degrade somewhat, but not seriously, as we showed in Section 4 (Figure 4.4). It is not unreasonable to hypothesize that a correct specification of the

[1]WSMIS is AFMC's Weapon System Management Information System. REALM (Requirements Execution Availability Logistics Module) is the part of WSMIS that computes WRSK requirements.

number of users in the spares requirements computation could overcome this degradation but, of course, the correct number of users is unknown without the configuration database.

The Air Mobility Command has already moved to a centrally managed WRSK, apparently very successfully, but it enjoys an essentially homogeneous configuration among its airlift aircraft of a given MDS. In the case of the Tactical Air Forces, when the need for deployment of a specified number of aircraft arose, a WRSK could be assembled and deployed to the desired location. Unfortunately, if there is significant heterogeneity among individual aircraft of a given type (model-design or MDS), there is no way for the central manager to know what specific stock numbers to include in the WRSK without an aircraft configuration database and a list of the tail numbers to be deployed.

Figure 5.3 reflects the investments associated with enough war readiness spares to support ODS and one-and-one-half and two times the ODS deployment for the F-16 aircraft. It also shows the investment required under the current policy of equipping every deployment-tasked squadron with its own WRSK as was done for the NATO scenario. A consolidated WRSK policy could, under current views about wartime contingency planning, deliver very attractive savings while not degrading *needed combat capability*. Moreover, its implementation could be coupled with a prudent WRSK withdrawal policy that would allow the withdrawal of some assets from the WRSK.

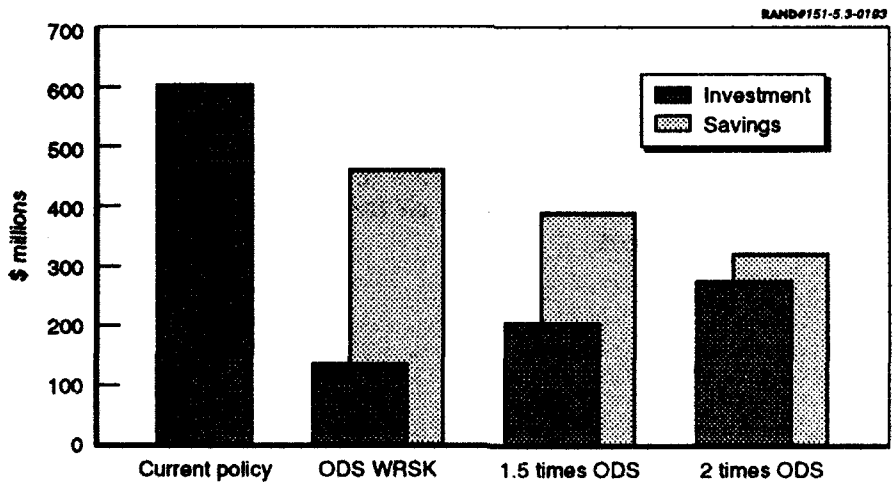

Figure 5.3—Savings Achievable with Alternative WRSK Policies

SRU SAFETY STOCK AT THE DEPOT

The current requirements computation allocates no SRU safety stock to protect the LRU depot repair pipeline. This implies that the LRU can wait for repair of needed SRUs, an imprudent assumption at best. It is a simple matter to provide for SRU safety stock at the depot, and it would pay off in enhanced robustness of depot repair. It is especially important for those LRUs whose repair process is such that lack of a serviceable SRU when needed induces inefficiencies. An excellent example of such a case is the repair of avionics LRUs on program-driven, automated test stations. The LRU is installed on the test station and the fault-diagnostic program checks the functions of the LRU until it either finishes the checkout or it discovers a defect of some sort. Typically, the identification of a defect requires the replacement of an SRU. If no serviceable SRU is available, the LRU must be removed from the test station and stored until an SRU becomes available. When the SRU is finally available, the LRU must be removed from storage, reinstalled on the test station, and its checkout started over again. Several hours can be lost in this procedure, to say nothing of the elapsed time the expensive LRU waits for the SRU to become available. Our experience at the Ogden ALC in the development and demonstration of the DRIVE prototype suggests that lack of serviceable repair parts for many repair processes is very costly in both LRU waiting time and total shop throughput capacity.

Evaluation of this policy recommendation is beyond the scope of this work. Its difficulty lies in collecting sufficient data to estimate awaiting parts (AWP) delay times for the current and proposed policies. Nevertheless, SRU safety stock at the depot is a concept we recommend for immediate implementation. Cost is obviously a consideration here, but our intuition suggests that, at least for avionics component repair and repair of other assets whose repair processes share the characteristic of being vulnerable to lack of serviceable repair parts, SRU safety stock would pay off far more in depot responsiveness than it would cost.

6. ESTIMATING DEPOT COMPONENT REPAIR REQUIREMENTS

Abell et al. (1992) analyzed the application of DRIVE to the estimation of quarterly repair requirements. DRIVE is a computer-based optimization algorithm that prioritizes depot-level component repairs and asset allocations to achieve specified aircraft availability goals. It takes quite a different view of the management of depot-level component repair than does AFMC's traditional component repair management system, called MISTR (Management of Items Subject to Repair). For a more complete discussion of DRIVE's application to the management of depot-level component repair, we refer the reader to Abell et al. DRIVE was demonstrated in prototype form at the Ogden ALC where it was applied to the prioritization of F-16A/B avionics LRUs and SRUs and their allocation to bases worldwide. Miller and Abell (1992) describe the details of this prototype.

The problem of estimating repair requirements is logically straightforward if the objective function is specified. The specification of a set of aircraft availability goals by model-design aircraft and location implicitly defines a *desired asset position,* i.e., a mix of serviceable spares at each location that will achieve a specified probability of meeting all of the aircraft availability goals. The depot repair system simply needs to repair assets to:

- Achieve the desired asset position (the *catch-up* requirement), and
- Maintain it (the *keep-up* requirement).

The total repair requirement is simply the algebraic sum of the catch-up and keep-up requirements. Obviously, the keep-up requirement can only be estimated at the start of the quarter; it is a random variable, the number of NRTS actions that will occur during the quarter. Clearly, the catch-up requirement could be negative, i.e., the asset position at the start of the quarter could be sufficiently rich that it could be allowed to degrade somewhat during the quarter and still achieve the desired system performance. Note that the catch-up requirement, too, can change during the quarter as the result of changing availability goals, revised estimates of item characteristics such as demand rates, changed missions for particular units, etc., further reinforcing the need to be adaptive in component repair management. As we have suggested, the current system doesn't compute the catch-up re-

quirement very well because of its lack of visibility of the evolution of the asset position and the force beddown; however, it does estimate the keep-up requirement in a logically straightforward way.

RAND's past research demonstrated DRIVE's effectiveness in estimating quarterly component repair requirements. It was shown to be clearly superior to the traditional approach for reasons that are discussed at length in Abell et al. (1992). Thus we have no doubt that a DRIVE-like mechanism should be the computing engine at the heart of whatever depot component repair planning system is eventually implemented by AFMC.

Ongoing research at RAND will extend the application of the logic underlying DRIVE to a decision support system for depot repair planning that is intended to enable decisionmakers to explore budgetary requirements and allocations, capital investment requirements, the effects of repair capacity constraints, allocations of assets to contract repair, the implications of repair mix for stock fund operating position and system performance, and other issues. It is also intended to explicate the keep-up requirement and illuminate decisions about allocation of the catch-up requirement over quarterly, annual, and longer planning horizons.

Since the demonstration of the DRIVE prototype at Ogden, AFMC has moved toward implementation of a production version of the system. We have several suggestions about the usefulness of that system in estimating quarterly component repair requirements. They involve: (a) data and data-processing issues, and (b) the applicability of prioritization logic to specific shops and workloads.

DATA AND DATA-PROCESSING ISSUES

We have examined the application data from which component indenture relationships are inferred and have concluded that a massive, intensive, continuing training program is needed throughout AFMC for IMs, equipment specialists, and others who define the data input to the D041 database. There are many troublesome issues associated with the database. Abell and Finnegan (forthcoming) discuss those issues in considerable detail. For repair requirements estimation, the issues are similar to those affecting spares requirements, except that DRIVE is even more sensitive to misidentification of SRUs than is the spares requirements computation. DRIVE gains considerable leverage over the prioritization and allocation problem through the rather surgical allocation of SRUs to alleviate AWP shortages, thereby generating additional serviceable LRUs through SRU allocations. That

leverage is lost if the indenture relationships among components are not correctly identified.

The database used to support the demonstration of the DRIVE prototype was essentially hand-built. The indenture relationships were, of course, well known to maintenance technicians who were intimately familiar with the components involved in the demonstration. It is quite a different matter to rely on a database that is fraught with errors that result from misidentification of component indenture relationships.

The principal problem in identifying SRUs and tying them correctly to their LRU parents is that, when the LRU record that should specify the linkage between the SRU and the aircraft is missing or reflects a different NSN than the SRU application, the SRU is promoted to an LRU. While such a scheme might be sensible for a small proportion of SRUs in the system in the context of spares requirements estimation, it is considerably more self-defeating in DRIVE's prioritization and allocation decisions.

We recommend implementation of the training program that we mentioned previously for people involved with the D041 database, especially the application file, at least for the purpose of supporting DRIVE. The rules for making entries in the database must be very clearly thought through and specified, and personal attention must be given to every person who is involved in the act of supplying or maintaining data.

Another troublesome issue with the database is the accurate estimation of SRU replacement factors. SRU consumption data are needed for each LRU type, and a way must be found to acquire them. There is good reason to believe that the SRU replacement factors contained in D041 seriously underestimate SRU replacements owing to the rather widespread use of wash-post procedures in SRU repair.[1] Wash-post actions are not reflected in the replacement factors.

The AWP data reflected in D165[2] are incomplete because no application is shown for many AWP shortages. This problem, too, inhibits the system's effectiveness.

---

[1]Wash-post procedures are transactions intended only to record the custodial responsibility for physical possession of an asset. They work very much like a system of hand receipts. They are not input to any data system.

[2]D165 is the AFMC's Mission Capability Requisition Status Reporting System. It reflects shortages of repair parts that induce repair delays of higher assemblies.

Other problems with the requirements database that are discussed by Abell and Finnegan (forthcoming) must also be explicitly considered in the context of repair requirements estimation and corrected. The confidence that people will have in the production version of DRIVE may otherwise be seriously undermined.

APPLICABILITY OF DRIVE TO SPECIFIC SHOPS AND WORKLOADS

Although DRIVE, coupled with an aircraft configuration and force beddown database, can be used effectively to allocate virtually any serviceable asset emerging from depot repair, its repair prioritization logic is not very useful for end items with long repair times or items that are extensively job-routed; therefore, it should probably not be applied to all recoverable items undergoing depot repair. The planning horizon has to be tailored to the flow times of individual items and shops. Moreover, if repair resources are not shared by more than one item, prioritization may not provide significant gains in responsiveness over the traditional system; i.e., scope of repair is a necessary condition to achieve significant leverage through prioritization. (Obviously, DRIVE is always useful for quarterly planning and asset allocation decisionmaking.)

This suggests a detailed evaluation of each ALC to define the subset of shops or repair resources for which repair prioritization makes sense. Moreover, the constraining resource in the repair of each item or class of items must also be defined because DRIVE is an algorithm that solves an optimization problem with a single constraint, at least in its current form.[3] Maintenance technicians could also help verify or correct the indenture relationships in the D041 application file and the DRIVE database. Confidence in the indenture data could be maintained through a system of audits of LRU families that would ensure that the application file is being maintained correctly. Capacity constraints should also be examined at each ALC to ensure that they do not inhibit the execution of the repair requirements.

AN ANALYSIS OF DRIVE'S REPAIR REQUIREMENTS ESTIMATION

In the remainder of this section, we analyze DRIVE's effectiveness in repair requirements estimation in contrast to the MISTR system,

[3]Louis W. Miller and Kenneth J. Girardini of RAND have developed and successfully demonstrated an alternative version of DRIVE for the U.S. Army that will accept more than one constraint.

using some recent data from the Ogden ALC. This discussion presents the results of an analysis performed by Louis W. Miller of RAND. The situation at Ogden has evolved over time in a way that has resulted in a paucity of serviceable assets at F-16 bases for some of the items whose repair is prioritized by DRIVE. The issue seems to be one involving repair capacity. In this analysis, we will explicitly examine the differences in the quantities of repairs estimated by DRIVE and those estimated by item managers in conjunction with production management specialists, which we refer to as the X21 scrubbed requirement.[4] The details of our analysis would probably fill about 30 pages; therefore, we have chosen one of the four test stations used to diagnose discrepancies in F-16 avionics components at Ogden, the computer-inertial (CI) test station, for illustrative purposes. The results for the other three test stations may be found in the appendix, but the CI example illustrates most of the kinds of differences we observed between DRIVE and MISTR repair requirements.

DRIVE was originally designed to prioritize component repair; i.e., it only suggested what components to repair over a short production period to maximize the probability of meeting specified aircraft availability goals. DRIVE has been modified by Miller (Miller and Abell, 1992) to adapt it to estimate quarterly repair requirements. It does not explicitly address the allocation of repair capacity across the various workloads in the depot, but it does estimate the globally optimal allocation. It also illuminates the decision about repair mix among various LRU families, between LRUs and SRUs, and among work centers that are involved in the repair of the DRIVE items. The DRIVE Decision Support Program (DDSP), also described in Miller and Abell (1992), is very helpful in these allocation decisions.

Figure 6.1 reflects the IM scrubbed requirement for DRIVE items in standard hours (IM Total Standard Hours), the X21 scrubbed requirement, and the workload negotiated at Ogden for the first quarter of FY 1992 (1 October through 31 December 1991), and compares them to the keep-up requirement. The keep-up requirement is simply the expected number of NRTS actions during the quarter, translated to standard hours of repair. Note that both the X21 scrubbed requirement and the negotiated workload fall short of the keep-up requirement. This means that unless substantial numbers of compo-

[4]The X21 scrubbed requirement is the IM scrubbed requirement modified by the additional quarters of data contained in the X21 report; it is the repair quantity that constitutes the basis of the quarterly repair workload negotiations.

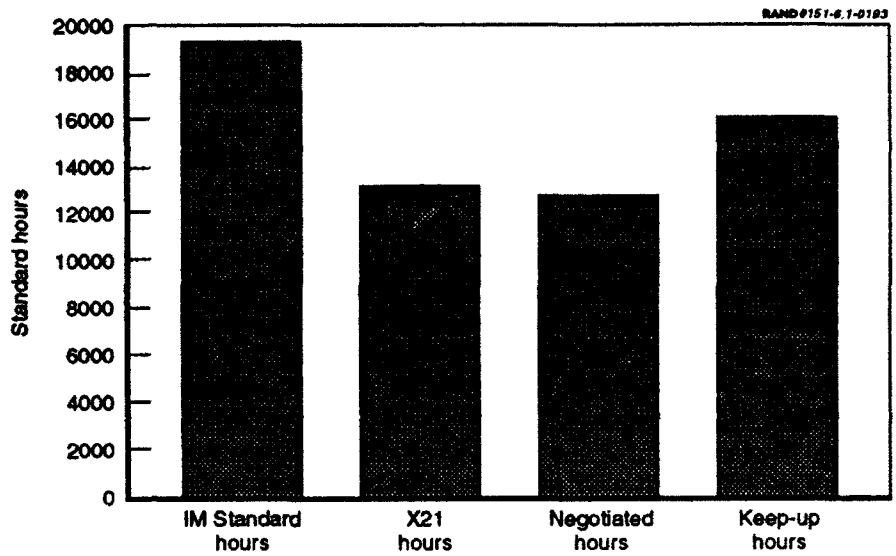

Figure 6.1—Total IM, X21, and Negotiated Requirements vs the Keep-Up Requirement

nent repairs are done on contract, the asset positions at the bases can be expected to deteriorate during the quarter. In addition to the negotiated number of repairs, it would require almost 4000 more hours just to avoid falling further behind, let alone catching up. This simple, straightforward calculation of the keep-up requirement is fundamental to repair requirements estimation; yet, to the best of our knowledge, it is not made visible in any of AFMC's standard data systems.

In the discussion that follows, we compare DRIVE's repair requirements for the CI test station with the X21 scrubbed requirements. To estimate repair requirements with DRIVE, we used the same total hours of capacity as represented by the total X21 scrubbed requirements for all the DRIVE LRUs. We ran DRIVE in a global optimization mode; i.e., we ignored capacities of individual test stations assuming that some repairs could be done contractually to relieve test-station-specific constraints. We then examine the asset positions of those items for which there was a significant difference between the DRIVE and X21 requirements in an effort to understand which of the two estimates is more closely tied to the actual asset position at the start of the quarter.

Figure 6.2 shows the X21 quarterly scrubbed requirement, DRIVE requirement, and keep-up requirement for each LRU repaired on the CI test station. The LRUs are identified by the last four digits of their stock numbers. The X21 requirement exceeds the DRIVE requirement for LRUs 3533 and 3160, but the opposite is true for LRU 6645. Since the DRIVE requirement is so much larger than the keep-up requirement for LRU 6645, it implies that the asset position suffers from a serious paucity of assets, relative to other LRUs. To understand which of these repair requirements estimation methods is more closely tied to the worldwide asset positions and the keep-up requirement, we show the asset positions of these three LRUs in Figures 6.3 through 6.5.

The asset position of the fire control/navigation panel (3533) is shown in Figure 6.3. The height of the black bar at each base reflects the total expected demand at that base across its entire planning horizon. The unit at Selfridge, the first base shown, has a wartime deployment mission. Thus the total expected demand of about 11 includes demands in both the peacetime and wartime horizon. The peacetime horizon is specified to be 18 days plus a base-specific order-and-ship time, 9 days in this case, for a total peacetime planning horizon of 27 days. Expected demands during the peacetime planning horizon are

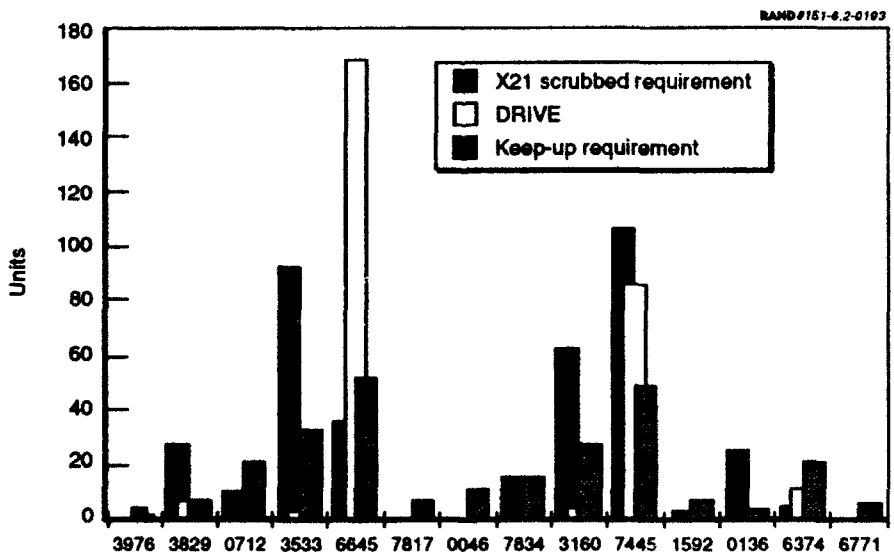

Figure 6.2—Repair Requirements for the CI Test Station

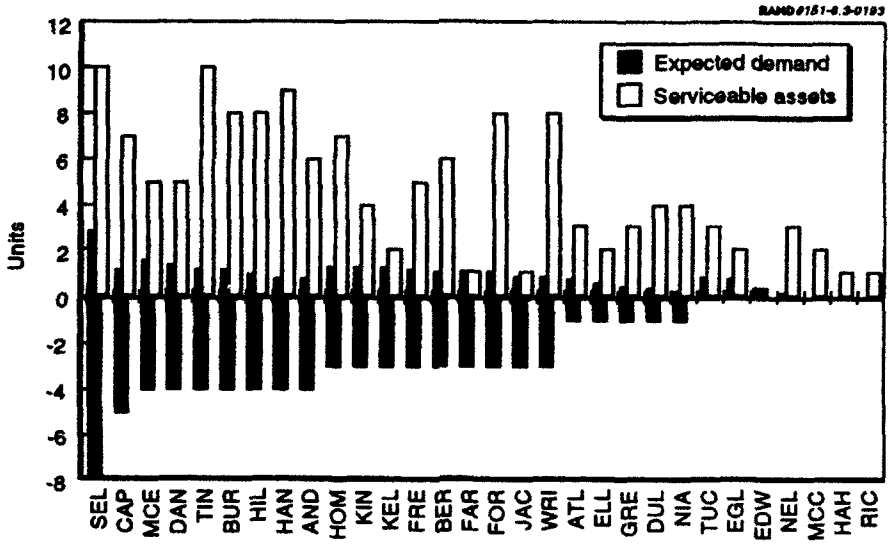

Figure 6.3—Asset Position of the Fire Control/Navigation Panel (3533)

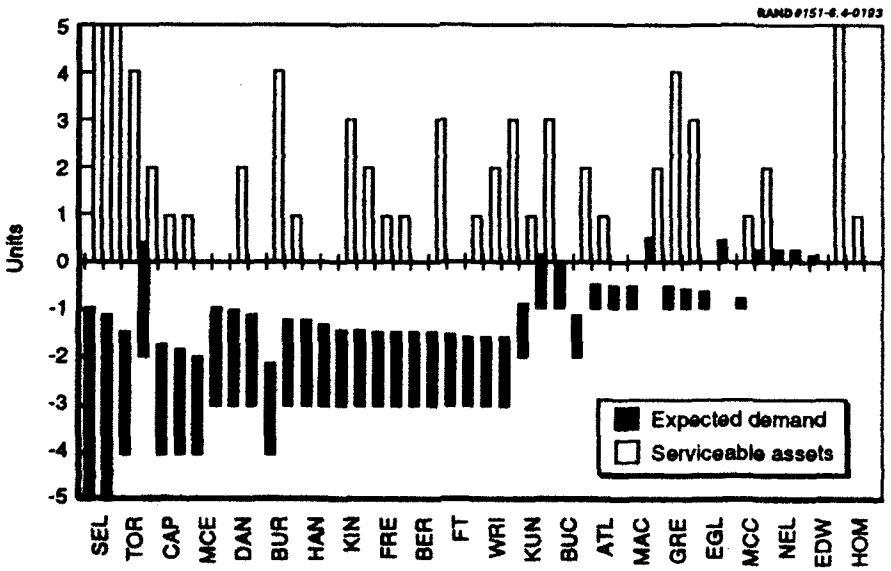

Figure 6.4—Asset Position of the 3160 Panel

estimated in a straightforward way by multiplying together the item's daily demand rate, quantity per application, application percentage, NRTS rate, and flying hours. This expectation is pooled with the expected demands in 30 days of wartime, which are computed using wartime flying hours and 100 percent NRTS rates.

The black bar is displaced below zero in this case to represent the number of allowable shortages of this particular component given the aircraft availability goal at the end of the total peacetime plus wartime planning horizon. With 54 aircraft and an 85 percent aircraft availability goal, Selfridge needs 46 aircraft available to meet the goal, i.e., it can have 8 aircraft unavailable at the end of the wartime planning horizon. Thus the bar representing expected demands is displaced downward by 8.0.

The asset position of the 3533 panel shows enough serviceable assets on hand to cover all the expected demands, a relatively rich asset position (although some lateral supply may be needed). Yet, the X21 requirement suggests far more repairs than needed to meet the keep-up requirement, implying the need to enhance the richness of an already rich asset position. Figure 6.4 reflects a similar situation for the 3160 panel assembly. Although its asset position is slightly less rich, enough assets are available to cover expected demands. Again, some redistribution actions may be needed.

On the other hand, the asset position of the inertial navigation unit (6645) shown in Figure 6.5 reflects a severe shortage of serviceable assets worldwide. There is virtually no chance of satisfying all of these demands with the assets on hand; nor is there much more chance to satisfy them with the number of repairs reflected by the X21 scrubbed requirement. The F-16 force was in very serious trouble on the inertial navigation unit at this time, but the X21 requirement is actually less than the keep-up requirement, implying that the asset position can be allowed to degrade during the quarter, an estimate of the repair requirement totally out of touch with reality.

We do not know why the current system frequently estimates repair requirements in ways that seem so inconsistent with the worldwide asset positions at the bases. We have observed this behavior in the past and are observing it again in this fiscal quarter. On the basis of our examination of data for all the test stations, we conclude that DRIVE makes much more sensible estimates of repair requirements than does the current system and that, given accurate data reflecting worldwide asset positions, item characteristics, and indenture relationships, DRIVE will yield a much more cost-effective mix of repair

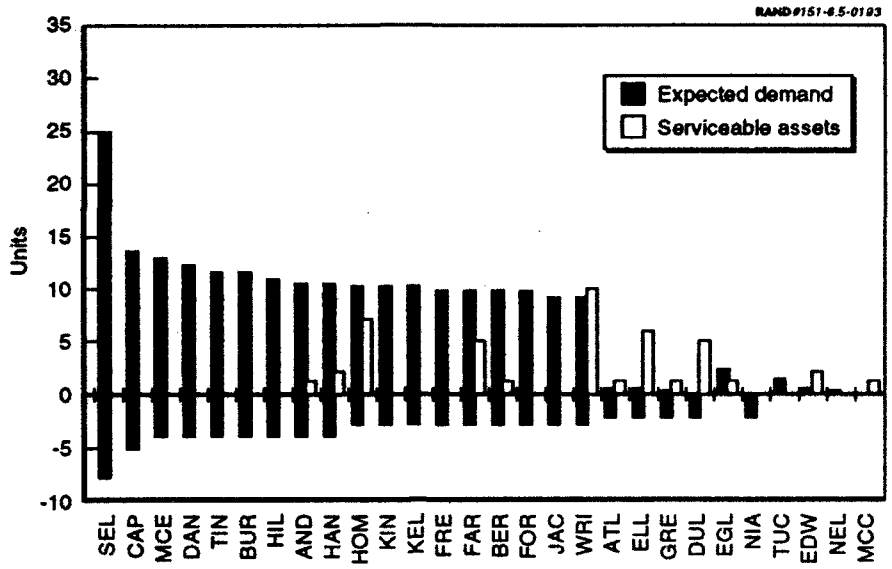

Figure 6.5—Asset Position of the Inertial Navigation Unit (6645)

actions to achieve specified aircraft availability goals (in terms of standard hours or other marginal cost) than will the current system. Note that we used the prototype version of DRIVE for this analysis. We have not examined the performance of the production version of the system in estimating repair requirements. We would expect the results to be approximately the same, though, except as they might be affected by data problems. The appendix contains data on the other test stations and the asset positions of those LRUs for which the X21 scrubbed requirement and DRIVE requirement differed substantially.

Estimating the Catch-Up Requirement

While the estimation of the keep-up requirement is straightforward, the estimation of the catch-up requirement is not. It requires definition of a *desired asset position* against which the catch-up requirement can be measured. The desired asset position is a function of the aircraft availability goals specified for each of the bases, the distribution of expected demands, the length of the planning horizon, and the assumptions we make about redistribution. Determination of the desired asset position also requires the use of rather sophisticated computational algorithms. DRIVE is imbedded in the software required to estimate quarterly repair requirements, of course, but additional procedures are required. Although most of these procedures are de-

scribed in Miller and Abell (1992), additional improvements in the software package have been implemented since then.

DRIVE operates on one LRU family (an LRU and all of its SRUs) at a time. One problem with estimating the catch-up requirement is how to decide when to stop repairing additional assets. We explore here, for the first time, a stopping criterion that we have found to be very useful in estimating the catch-up requirement.

One possible stopping criterion within an LRU family is the probability of achieving the specified aircraft availability goals at all bases. The problem with this criterion is that its numerical value depends heavily on the number of bases with positive expected demand. The larger the number of bases, the smaller the stopping criterion needs to be because the probability of meeting the availability goals at all the bases is simply the product of the probabilities of meeting the goal at each individual base. A criterion that is essentially immune to inter-item heterogeneity in the number of bases with positive expected demand is the geometric mean of the probabilities at the individual bases. The geometric mean is simply the nth root of the product of n numbers. Thus, if an item has positive expected demand at 27 bases, the geometric mean would simply be the 27th root of the product of the 27 probabilities. (It is also the antilog of the mean of the logarithms of the 27 probabilities.) This simple device enables us to equate performance goals for items that are applicable to a few bases with those for items that are applicable to many bases.

In the discussion that follows, we used this geometric mean to help us determine the catch-up requirement. We specified its value to be 0.95 for all LRU families, then we ran DRIVE in two ways: (a) taking explicit account of carcass constraints, and (b) ignoring carcass constraints. In the former case, we stopped repairing when we reached the 0.95 goal or when we ran out of carcasses, whichever occurred first. The results are shown in Figure 6.6. The keep-up requirement is unvarying, of course. It is simply the number of standard hours of repair needed to keep up with the repair of carcasses as they generate. DRIVE hours (U) represents the total number of standard hours of repair required to achieve the 0.95 goal for every LRU family, ignoring carcass constraints. It is the sum of the keep-up and catch-up (U) requirements. Without sufficient carcasses to reach the 0.95 goal, we are constrained to a catch-up ability of less than 10,000 standard hours. No carcass constraint applies to the keep-up requirement, of course; the carcasses are generated through the NRTS actions of the bases. Only the catch-up requirement needs to be covered by repairable carcasses at the start of the quarter.

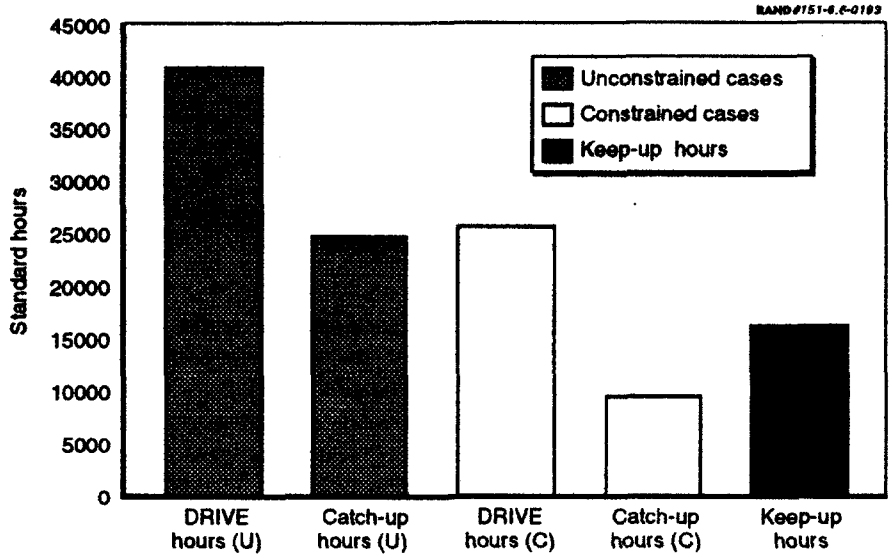

Figure 6.6— Estimation of the Catch-Up Requirement

Beyond carcass constraints, the DRIVE prototype has three additional ways of enabling a user to specify a stopping criterion within an LRU family: (a) the probability of meeting the aircraft availability goals at all bases, (b) the sort value, i.e., a measure of the improvement in the probability of meeting the goals divided by the unit repair cost, and (c) the repair capacity constraint. These several criteria can be used singly or in combination. The geometric mean is an improvement to the probability of meeting the availability goals at all bases. Its use differs from that of the sort value, which enables one to find an "optimal" stopping point for each LRU family that equates to a line drawn on a globally optimized repair list. The geometric mean does not account for repair costs; the sort value does.[5] Thus, the geometric mean does not yield a solution that is optimal in the sense of getting the most good out of limited repair resources. It does, on the other hand, provide a measure of performance that is intuitively more sensible than the probability of meeting aircraft availability goals at a large number of bases.

It is important to note that the version of DRIVE we used here incorporates a redistribution assumption. During demonstration of the

[5]Additional research and evaluation is needed to determine how important it is to account for repair costs in moving to the use of a geometric mean.

DRIVE prototype at the Ogden ALC in 1987, we performed a similar analysis of quarterly repair requirements and reached the same conclusions, although without the redistribution assumption. The earlier analysis is described in Abell et al. (1992), in which the redistribution issue is discussed in depth. The data in both of these analyses are compelling and mutually reinforcing. They both conclude that a DRIVE-like mechanism for estimating quarterly repair requirements is clearly superior to the current system. Coupled with a resolution of the data issues that trouble AFMC's standard data systems, such a mechanism, as the computing engine in whatever repair planning and management system ultimately evolves in AFMC, could be a powerful aid in decisionmaking about depot-level repair.

DRIVE AND THE CRITICAL ITEM PROGRAM

Another question of interest in this analysis was whether there was much consistency between the items whose repair DRIVE was most highly prioritizing and the items that Headquarters, TAC, judged to be "critical" at the time. To evaluate the level of consistency, we built a "critical item list" with DRIVE using the three criteria listed below and compared that list of items with the TAC critical item list. The DRIVE database used in this analysis was dated 30 September 1991; the TAC critical item list was dated 1 October 1991. The criteria we used were:

- The item's catch-up requirement exceeded its keep-up requirement (in other words, the depot was more than a full quarter behind in repair).
- The item's catch-up requirement exceeded 500 hours.
- No matter what repair capacity was allocated to the item, its geometric mean would not reach 80 percent before the supply of repairable carcasses was exhausted.

Note that the first and second of the above criteria are readiness related because they measure the adequacy of the number of serviceable assets in the system, while the third is really a measure of sustainability in the sense that it expresses a paucity of total assets in the system. Thus, either readiness or sustainability deficiencies can cause an item to be on DRIVE's critical item list, just as these deficiencies can cause items to be on the TAC list. The items on the DRIVE critical item list are identified by an *a*, *b*, or *c* if they met criterion one, two, or three, respectively. The list is shown in Table 6.1.

Table 6.1

Comparison of the TAC Critical Item List with a DRIVE List

TAC Critical Item List	DRIVE Critical Item List
CI6645 INU	CI6645 INU[a,b,c]
DI4833 REOEU (74EB0)	DI4833 REOEU (74EB0)[a,b]
PP6879 SCP	DI7430 HUD PDU[a,b]
DI7430 HUD PDU	RF296X LPRFs[a,b,c]
CI1859 CADC	PP1499 MRIU[a,b]
CI1018 CADC	PP6879 SCP[a,b]
	PP4855 CIU[a,c]

[a]Catch-up significantly exceeds keep-up.
[b]Catch-up exceeds 500 standard hours.
[c]Carcass-constrained below 0.80 geometric mean.

Note that the first four items on the TAC list also appear on the DRIVE list. The low power radio frequency (LPRF) unit on the DRIVE list also appears on TAC's "Problem Item List," a list of items that are causing some problems but are judged to be not as seriously in trouble as are the critical items. The missile release interface unit (MRIU) and central interface unit (CIU) on the DRIVE list have SRUs on the TAC critical item list. We judged the two central air data computers (CADCs) not to be as critical as TAC did; however, there may be issues involving modifications to these two LRUs of which we are unaware. Clearly, the two lists are very nicely correlated.

7. TRADEOFFS AMONG SPARES, REPAIR, AND TRANSPORTATION

In this section, we depart from the micro-level view of previous sections and discuss the possibilities for system-level tradeoffs among logistics resources. Many logisticians have long believed that significant gains could be made in cost-effectiveness if it were possible to reduce spares investment requirements by reducing resupply times. The Air Force's planned implementation of stock funding of recoverables may make such tradeoffs less difficult than they have traditionally been.

The characteristic of the current logistics management system that has historically inhibited such tradeoffs is that resource allocation decisionmaking is partitioned. Within each functional area, decisionmakers try to achieve cost-effective allocations without the cross-functional visibility of the costs and payoffs of alternatives. Moreover, resource allocations are shaped by decisionmaking at more than one echelon and are affected by policies formulated at various echelons. Thus, AFMC may decide through careful evaluation that it can make significantly more cost-effective expenditures if it shortens depot repair turnaround times by maintaining more repair capacity and better responsiveness than it has had in the past and takes the gains in reduced spares investments. It may decide, too, that it needs more responsive transportation for many items and richer stocks of consumable and recoverable repair parts in the depot materiel support system than current policy provides.

Suppose the increased expenditures to achieve this enhanced responsiveness resulted in savings in recoverable spares five times as large, let us say, as the additional expenditures required to achieve the higher levels of responsiveness. On the face of it, the choice may be obvious but, as a practical matter, strategies that appear optimal from a systemwide perspective often appear suboptimal from a functional area perspective and may be difficult to implement because of the partitioning of decisionmaking and responsibilities in the system. Virtually every interested agency in the Air Force, OSD, and the Congress would need to understand the several ingredients of the strategy. For example, such an approach could decrease use of repair resources in the depot system. Depots might be judged to have "excess" capacity, insufficient turnover in their inventories of repair parts, or underused maintenance technicians. Thus a strategy that is

judged to be cost-effective from a system point of view can have serious implications for managers throughout the system as a result of the partitioning problem.

Evaluation of what we called the quick response option in Section 4 has strong implications for resource allocation tradeoffs at the system level. Modest reductions in depot repair times can yield very significant savings in spares investment requirements. The distribution of depot repair times for active BP-15-coded items reflected in the D041 database is shown in Table 7.1. These times include both depot repair flow time and retrograde time, i.e., they represent the elapsed time from removal of a component from an aircraft until it is serviceable in depot supply, but they are defined to exclude waiting time in repairable storage. This distribution is characterized by a very long positive tail that extends all the way to 450 days, although only 5 items in 100 have depot repair times above 127 days. If the items with very long depot repair times are not items whose availability seriously affects system performance, their depot repair times may not be of serious concern.

Consider the effects on spares requirements estimation of implementing the quick response option that we used in the evaluations of Section 4. This would involve a reduction of total depot turnaround time (in this case including order-and-ship time) from an average of 89 days to about 50 days. The quick response option did not assume any reduction in shop flow days, only in transportation and handling. With such reductions, and prioritized repair of the most urgently needed items, item pipelines could be reduced by roughly three-fourths of a billion dollars *with no other initiative.*

Table 7.1

Distribution of Depot Repair Times

Days	% of Times ≤ This Number
10	1.0
30	9.7
45	49.5
60	71.3
75	85.0
90	89.4
105	91.1
120	92.8
135	95.8

8. CONCLUSIONS AND RECOMMENDATIONS

After extensive analysis and evaluation, we come away impressed with the difficulty of estimating spares and repair requirements. The system is large and complicated, spread across five air logistics centers in execution, and fraught with uncertainty, the vagaries of churn, and many other factors that are either invisible or unknown to us. Thus we need to qualify many of our observations, conclusions, and recommendations with caveats. We are constrained in our analyses by the quality of the data at our disposal, just as the requirements system is. Moreover, we have no real intuition about the magnitude of the difference between the mix of spares that is computed in the requirements process and the mix of spares that is bought, or at least specified in purchase requests by item managers. If the process of translating the computed mix of spares requirements into actual procurement actions is anything like the process by which item managers and production management specialists estimate repair requirements, there is good reason to believe that we are buying mixes of spares that are not only suboptimal, they are grossly suboptimal.

In general, we are persuaded that the system buys too many spares, and that the mix of those spares could be more cost-effective. But we fuss about such things as the specification of the number of users of a stock number, lack of consideration of base-specific application percentages and actual force beddown, multiplying too many factors together to estimate aircraft availability, exclusion of BCS demands from the computation of item pipelines, omission of engines and engine modules from the database, the substantial differences between the item programs in D041 and those implicit in the K004 and application files, SRU promotions to LRU status, and the like. Yet we are unable to evaluate to our own satisfaction how important any one of these apparent logical inconsistencies with reality really is. We can only point them out. Thus, after many months of investigation, analysis, and thoughtful reflection on the requirements problem, we still have many uncertainties of our own and have become truly respectful of the difficulty of the problem; however, we are persuaded about the payoffs of several recommendations that we make in the remainder of this section. We judge that, in total, they will result in substantial aircraft recoverable spares cost reductions and better system performance for the dollar.

We also feel obliged to point out that the current system excludes late deliveries of assets from procurement, awaiting parts, and awaiting maintenance times, disallows any variance-to-mean ratios greater than 5.0, and understates the number of users. Thus, we do not intend to imply that there is some sort of conspiracy operating here. The estimation of spares and repair requirements is simply a difficult problem for all of the reasons we have already pointed out.

Given all of this, we believe that the people involved in the requirements process should do several things, and evaluate several other things with further research, perhaps with their own resources, perhaps with contractual resources. Thus what follows is a combination of initiatives that seem to us to be clearly beneficial and ideas that need further evaluation. Such evaluation, however, requires aircraft configuration data and other data that are unavailable to us. In some instances, our evaluative tools fail us as well. Suffice it to say that this research is imperfect, in many ways incomplete, but some ideas have emerged that we judge to be worth implementing or pursuing further to enhance the rationality of the spares and repair requirements process and enhance it in ways that will deliver traditional levels of performance at reduced cost.

RECOMMENDATIONS

We recommend that AFMC:

- Implement a massive, intensive effort to correct the D041 application file, and a continuing training program in principles and procedures for specifying elements of data that describe item characteristics and applications and in file maintenance. This training program should be coupled with a system of audits of LRU families that are sampled from the population in a way that accounts for item managers, equipment specialists, and commodity groups. The audits should be done to find data problems and identify item managers and equipment specialists who need assistance or training. Results should be analyzed to identify systemic problems.
- Establish an aggressive and continuing program to enhance the responsiveness of depot component repair and the management of all segments of the depot repair pipeline, and couple those functions more closely to the combat force. Cost analysis should be a first step.
- Implement the improved demand forecasting and variance specification techniques described in Adams et al. (1993) as the default

methods for high-demand items and reduce investment levels to maintain desired system performance.

- Enhance the Aircraft Availability Model as described in Gaver et al. (1993) to support the policy of designated cannibalization and modify the requirements database to reflect the designation of selected items as cannibalizable.
- Add engines and engine modules to the requirements database.
- Develop a configuration database for all aircraft that contains a mapping of subgroup master stock numbers to aircraft serial numbers and a system for routinely updating it to reflect changes to the mapping.
- Provide for SRU safety stocks for the depot repair system in the spares requirements computation. Again, costs should be estimated as a first step.

The suggestions that follow depend on the aircraft configuration database being in place:

- Consolidate the storage and management of war readiness spares, either at a single location, a location for each command, or a location for each weapon system, and develop the ability to compute, assemble, and deploy war readiness spares kits in a very short time (say 24 hours), on very short notice (say 6 hours). This capability should include mission support kits, en route support kits, and similar contingency requirements.
- Compute quarterly component repair requirements using DRIVE as the computing engine. Extend such requirements computation over longer planning horizons as needed for resource allocation and capital investment decisionmaking.
- Evaluate the cost-effectiveness of base-specific data coupled with a model of lateral supply in the requirements computation. By "base-specific data" we mean the actual number of bases at which items are exposed to OIM demand rather than the *number of users* in the current system, and base-specific application percentages, expected demands, aircraft availability goals, and stock-level allocations.

The effects of modeling lateral supply and using base-specific data tend to be offsetting. Combining them, we believe, would yield a more effective mix of spares. The use of base-specific data would make the calculations of the AAM considerably more tedious unless the estimation of depot stock levels were partitioned from that of base stock

levels. Modeling lateral supply could also help alleviate the problem, depending on how it was done.

Lest we be regarded as somewhat cavalier in our recommendations, we hasten to add again that the evaluation of the costs of these initiatives is beyond the scope of this research. We realize, however, that data collection and maintenance are not free, especially the data needed to support consolidated WRSK and designated cannibalization. On the other hand, given that each of these initiatives could reduce requirements by hundreds of millions of dollars, it seems to us that such data collection is more than worth its cost.

Our remaining suggestions are more technical in nature. Although they may not yield large savings, they are important, and may be very important, to the effectiveness of the mix of spares that emerges from the requirements computation:

- Undertake further analysis to resolve the major differences between D041 item programs and those implied by the K004 and application files.
- Evaluate the cost-effectiveness of adding BCS demands to item pipelines.
- Delay the specification of the procurement quantities of spares being purchased as long as possible before obligating the government to the contract. Use this additional time to observe additional demand and NRTS data to revise the procurement quantity. (We believe that this initiative is already in place.) Couple this strategy with specification of procurement quantities that leave some or all of additional safety stock requirements as an option for future procurement.

We believe that these several suggestions would result in more effective mixes of aircraft recoverable spares being procured as well as substantial reductions in spares investments. Implementing improved demand forecasting and variance specification, reducing the investment level to maintain desired system performance, consolidating WRSK at, say, twice the size needed to support Operation Desert Shield/Storm, and implementing the designated cannibalization concept would have reduced the stock levels that emerged from the March 1986 computation by roughly a billion dollars. We do not understand the implications of implementing all of the suggestions described in this section together. Thus we have tried to distinguish ideas that we think should be implemented without further evalua-

tion from those whose effects we can only estimate given the limitations of our models and data.

We realize, too, that we are being somewhat imprecise about an implementation strategy, but our priorities are implicit in the list of recommendations, except for the last three, which can be undertaken immediately.

In summary, however, we are persuaded that the Air Force can achieve roughly its traditional levels of system performance at substantially reduced spares investment costs.

Appendix

ADDITIONAL DATA ON REPAIR REQUIREMENTS

We discuss here additional results of our analysis of repair requirements estimations done with the 30 September 1991 prototype DRIVE database. In Section 6 we discussed the results for the CI test station. In the discussion that follows, we will present our results for the other three test stations: RF (radio frequency), DI (display-indicator), and PP (pneumatic-processor).

THE RF TEST STATION

Figure A.1 portrays the X21 scrubbed requirement, DRIVE requirement, and keep-up requirement for each LRU diagnosed and repaired on the RF test station. Each F-16A/B aircraft is equipped with a low power radio frequency unit (LPRF). Four different LPRFs are stocked in the inventory system. They differ only in a single SRU that determines their operating frequencies, but each type LPRF has

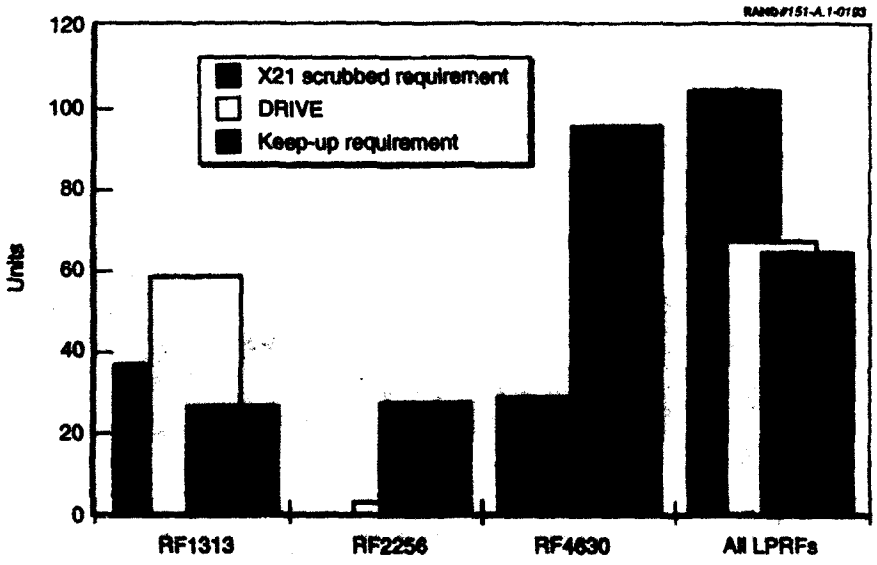

Figure A.1—Repair Requirements for the RF Test Station

a unique stock number. Data for the four types of LPRF are pooled in the rightmost set of bars. DRIVE is asking for more repairs of the RF1313 and fewer repairs of the RF4630 and LPRF than reflected in the X21 scrubbed requirements. The asset positions of these three LRUs are shown in Figures A.2 through A.4.

As before, the downward displacement of the black bars in the figures representing asset positions of individual LRUs reflect the allowable number of shortages of the LRU allowed at each base at the end of its horizon. Since peacetime availability goals are all specified to DRIVE to be 100 percent, but the availability goals specified at the end of the wartime planning horizon are only 85 percent, the downward displacement occurs only at bases with wartime deployment tasking.

A good way to gain some intuition about these repair requirements is to subtract the keep-up requirement from either of the estimated quarterly repair requirements, X21 or DRIVE, and add the difference to the current asset position. In Figure A.4, for example, the X21 repair requirement would result in about 39 additional serviceable LPRFs at the bases, while the DRIVE requirement would result in only 3 additional serviceables. The X21 requirement is 101. The keep-up requirement is 62.06. The difference, 38.94, or 39, is an es-

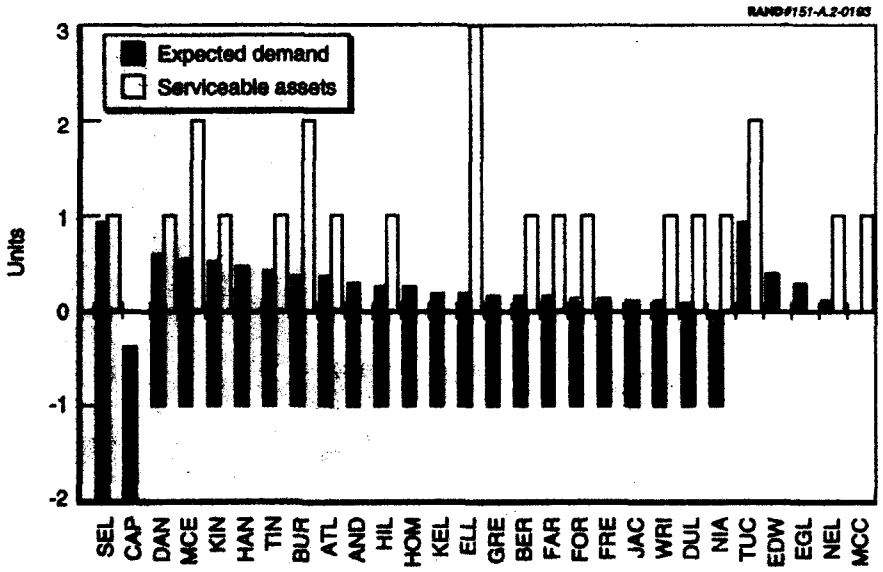

Figure A.2—Asset Position of the 63BG0 Indicator (RF1313)

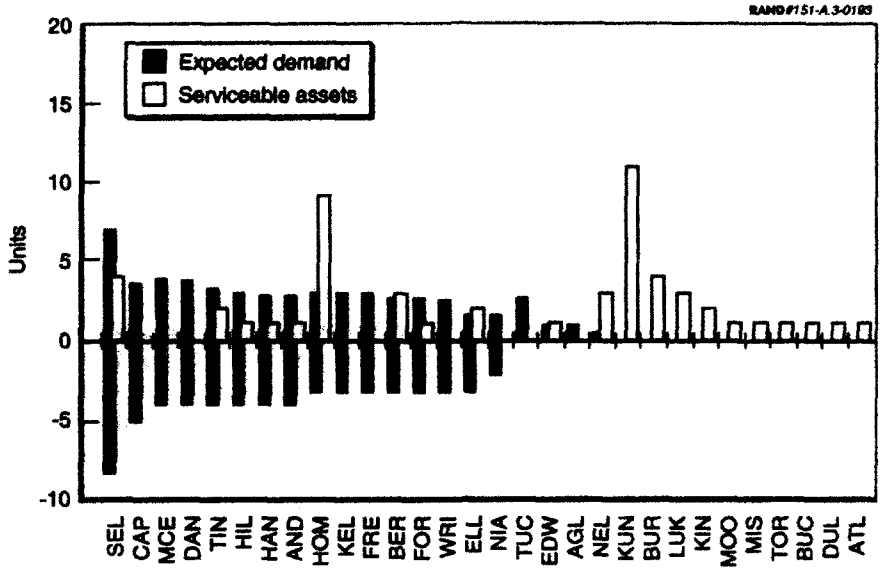

Figure A.3—Asset Position of the Radar Antenna (RF4630)

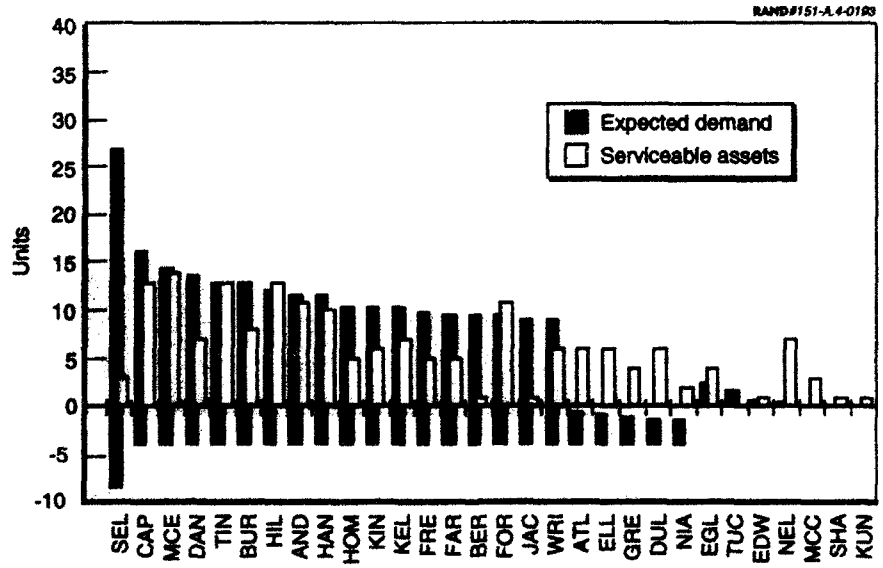

Figure A.4—Asset Position of the LPRFS (RF296X)

timate of the expected additional serviceable assets that would be gained from repairing the X21 scrubbed requirement. Clearly, 39 additional serviceable LPRFs would yield a much more comfortable asset position. The problem is that it would be relatively rich compared with other LRUs.

THE DI TEST STATION

Figure A.5 shows the results for the DI test station. Note that the X21 and DRIVE requirements are somewhat better aligned in this case; however, they differ quite dramatically in the case of the radar control panel (RF6872). DRIVE wants to add about 6 more LRUs to the asset position, while the X21 scrubbed requirement is satisfied to give up about 24, leaving only about 13 serviceables to cover more than 37 expected demands. Figure A.6 shows that DRIVE's results make more sense in the face of the asset position at the start of the quarter.

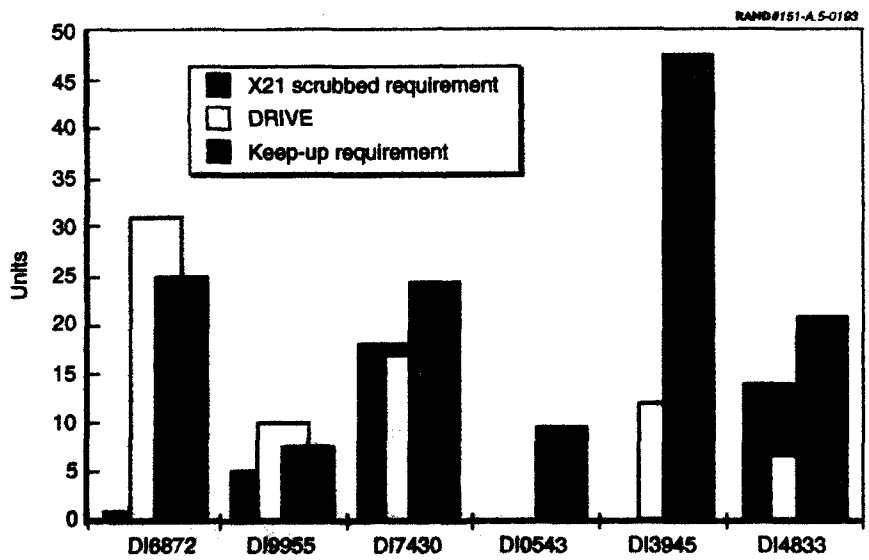

Figure A.5—Repair Requirements for the DI Test Station

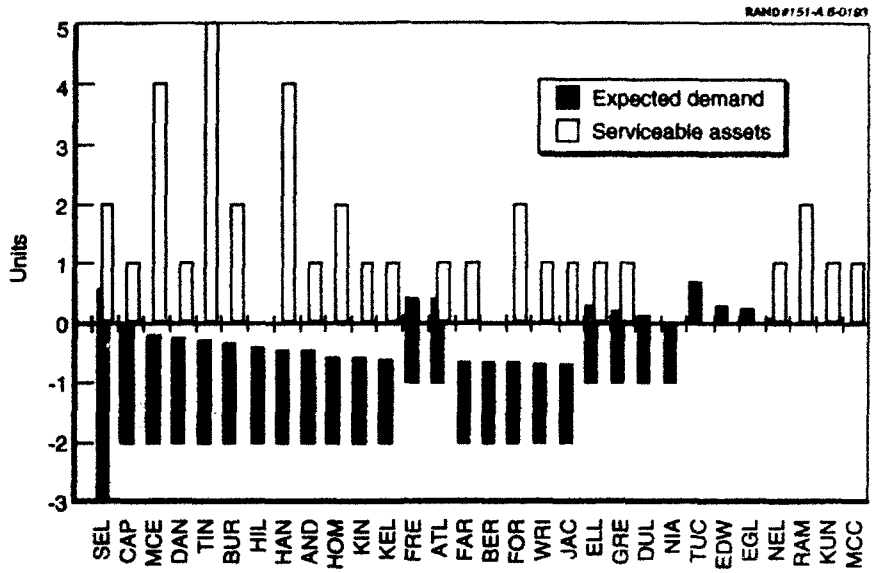

Figure A.6—Asset Position of the Radar Control Panel (DI6872)

THE PP TEST STATION

Figure A.7 shows the repair requirements for the PP test station. Note the striking differences between the X21 and DRIVE requirements here. We examine five of the asset positions underlying these repair requirements in the discussion that follows. We begin, in Figure A.8, with the LRU with the largest difference between the X21 and DRIVE requirements, the pneumatic sensor (PP6398). Proceeding as before, we note that the X21 scrubbed requirement would add about 113 serviceable assets to the asset position, while DRIVE would maintain the same number as are now available. Admittedly, this LRU has maldistribution problems. Some redistribution actions will no doubt be required. Nevertheless, the current asset position provides 28 serviceable assets to cover roughly 28 expected demands with 41 allowable shortages at the end of the horizon. The X21 repair requirement would provide 141 serviceable assets to cover the 28 expected demands, hardly a prudent action in the face of the asset positions of many of the other LRU families that share these test stations.

Figure A.9 portrays the asset position of the central air data computer (PP1859), for which the X21 repair requirement is zero, while the

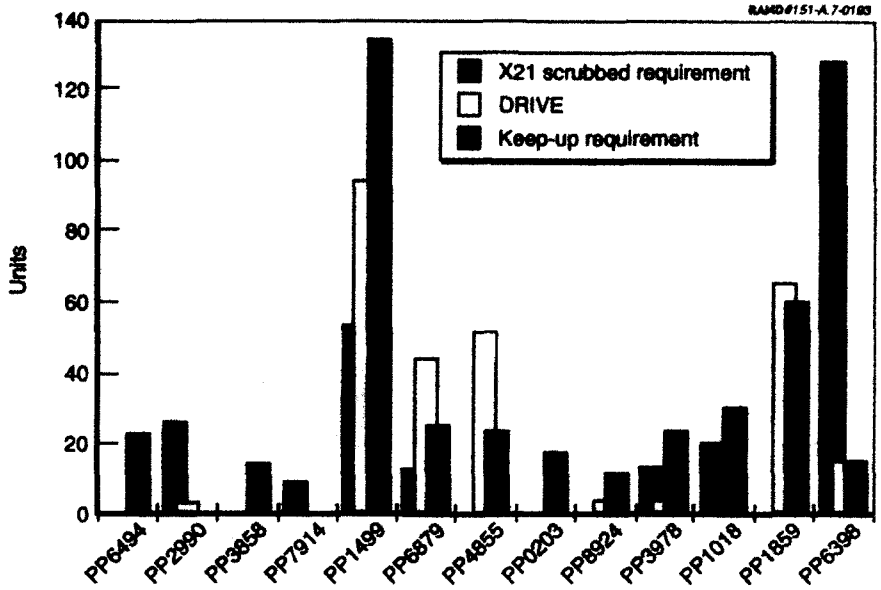

Figure A.7—Repair Requirements for the PP Test Station

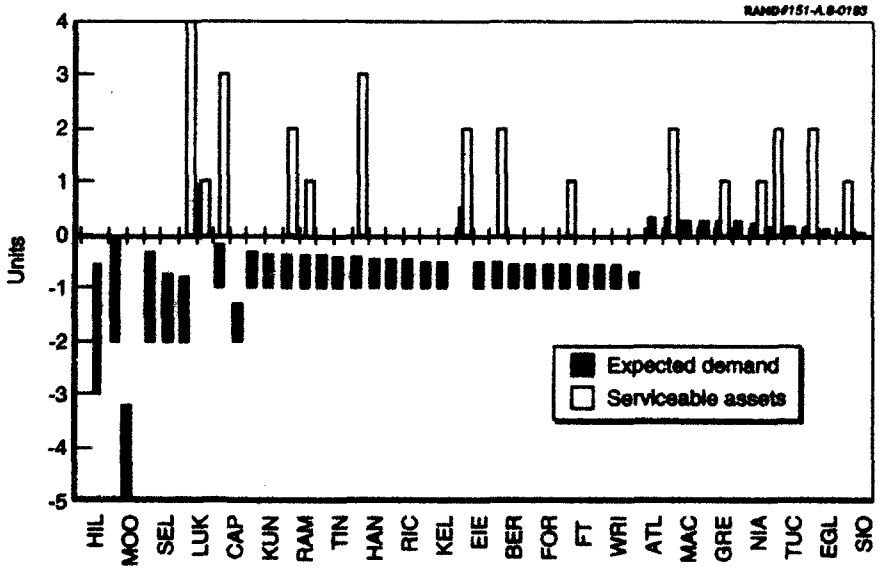

Figure A.8—Asset Position of the Pneumatic Sensor (PP6398 or 7835)

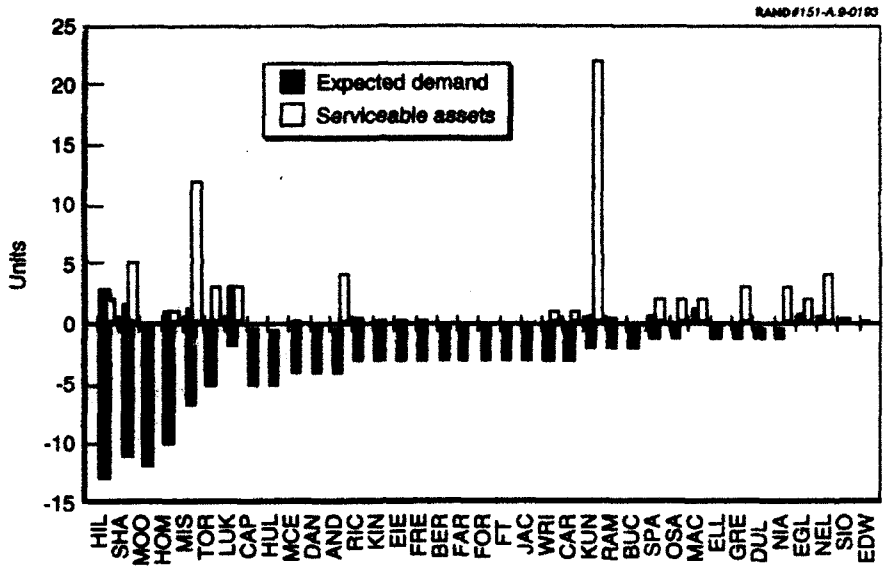

Figure A.9—Asset Position of the Central Air Data Computer (PP1859)

DRIVE requirement is 65, exceeding the keep-up requirement by about 5. There are only 72 assets of this LRU to cover about 146 expected demands; however, there are also 127 allowable shortages at the end of the planning horizon.

Note that there are two central air data computers (CADCs) for the F-16A/B aircraft; one is the PP1859 (6610-01-308-1859WF); the other is the PP1018 (6610-01-089-1018WF). The X21 scrubbed requirements allocate all repairs to the PP1018, while DRIVE allocates none, and vice versa; i.e., the X21 wants no PP1859 repaired, but allocates all repairs to the PP1018. We infer from this that the DRIVE database might not have been maintained properly to reflect the latest item configuration, the need for retrofit, or some other factor that is known to the item manager but not to DRIVE. All we can do here is show what makes sense given what is in the database. Clearly, if the database is not kept up to date, DRIVE's requirements may not be tied to reality. We are also able to conclude without doubt, however, that if the data are correct, DRIVE's repair requirements are much more sensibly tied to the asset positions than are the X21 requirements.

Figure A.10 portrays the asset position of the central interface unit, PP4855. There are only 14 assets of this unit to cover almost 90 ex-

pected demands with 24 allowable shortages at the end of the planning horizon. DRIVE is asking for 52 repairs, about 28 more than the keep-up requirement, while the X21 reflects no repair requirement. Again, we wonder about the quality of the database; given the data, however, Figure A.10 clearly indicates how disconnected the X21 requirement is from the asset position.

Two more examples should suffice for illustrative purposes: (a) the missile release interface unit (MRIU, PP1499), and (b) the stores control panel (SCP, PP6879). In the case of the MRIU, Figure A.11, the DRIVE requirement exceeds the X21 scrubbed requirement by 41 units. There are 221 assets to cover about 425 expected demands with 344 allowable shortages at the end of the planning horizon. The allowable shortages are so great because this unit has a quantity per application of 6. Neither the X21 nor DRIVE repair requirement will keep up with demands on this item; thus, its asset position will degrade during the quarter no matter which repair requirement applies. The DRIVE requirement will allow it to degrade by about 41 assets, the X21 scrubbed requirement by about 82 assets. Most of the shortages will be felt in war readiness spares.

The last asset position we examine is that of the stores control panel (PP6879) shown in Figure A.12. There are only 12 assets to cover

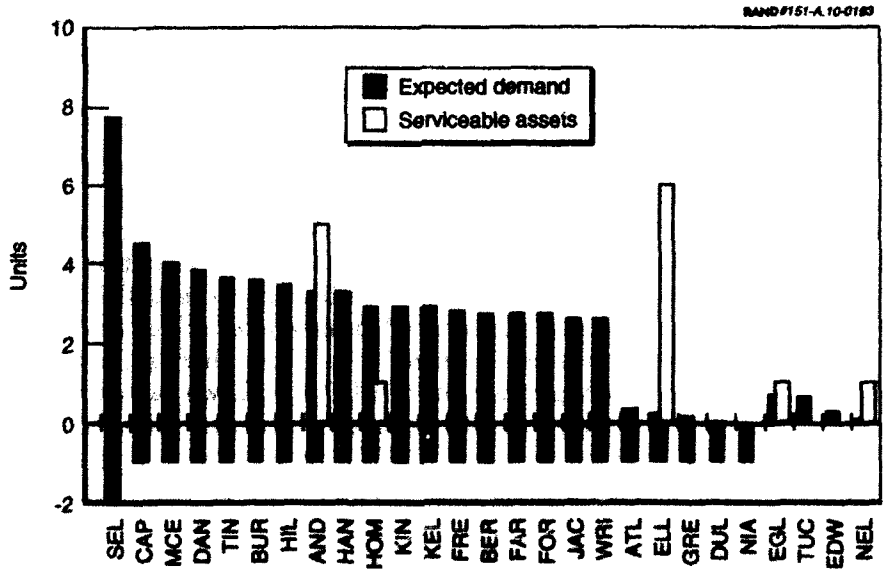

Figure A.10—Asset Position of the Central Interface Unit (PP4855)

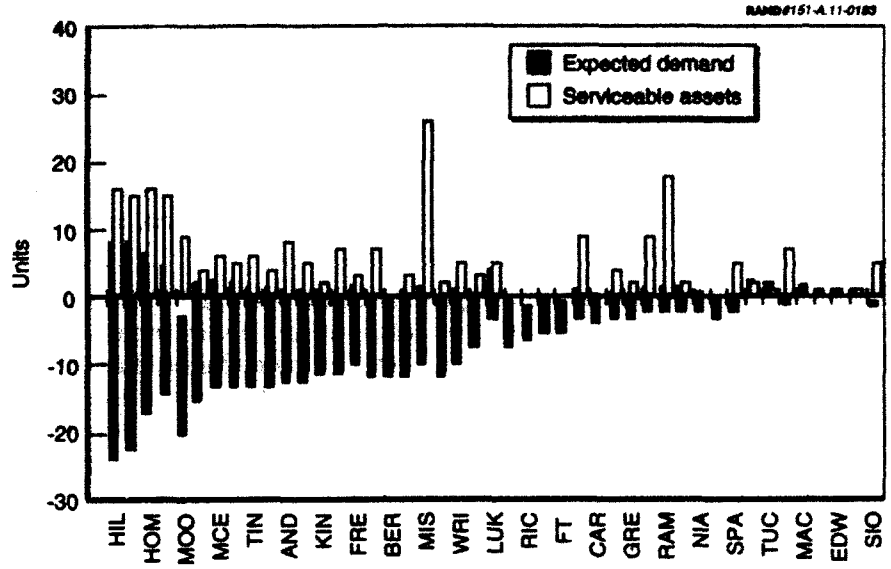

Figure A.11—Asset Position of the Missile Release Interface Unit (PP1499)

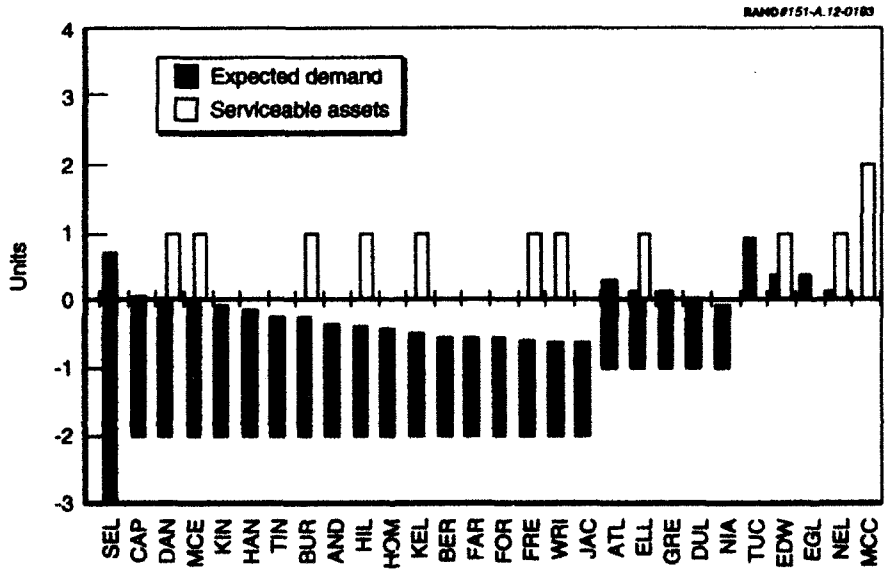

Figure A.12—Asset Position of the Stores Control Panel (PP6879)

about 40 expected demands with 42 allowable shortages at the end of the planning horizon. DRIVE is trying to improve the asset position by adding about 19 additional assets; the X21 scrubbed requirement would allow the asset position to degrade by about 13 assets, barely enough to meet the expected demands during the quarter. By implication, then, the X21 scrubbed requirement will allow all the panels at the bases to be consumed during the planning horizon, and shortages of about 13 panels will go unsatisfied.

In summary, we conclude that, given accurate data, DRIVE is clearly superior to the current system in estimating repair requirements. The problem with both the DRIVE prototype and the production version of DRIVE is data quality. This finding reinforces our recommendation for an intensive clean-up of the application file and other data used in spares and repair requirements determination, and a continuing program of training and data auditing.

BIBLIOGRAPHY

Abell, John B., and Frederick W. Finnegan, *Data and Data Processing Issues in the Estimation of Requirements for Aircraft Recoverable Spares and Depot Repair*, RAND, MR-264-AF (forthcoming).

Abell, John B., *Estimating Requirements for Aircraft Recoverable Spares and Depot Repair: Executive Summary*, RAND, R-4215-AF, 1993.

Abell, John B., et al., *DRIVE (Distribution and Repair in Variable Environments): Enhancing the Responsiveness of Depot Repair*, RAND, R-3888-AF, 1992.

Adams, John L., John B. Abell, and Karen E. Isaacson, *Modeling and Forecasting the Demand for Aircraft Recoverable Spare Parts*, RAND, R-4211-AF, 1993.

Air Force Logistics Command, *Recoverable Consumption Item Requirements System (D041): Studyguide/Workbook, Segments 0 thru 5*, December 1984.

Coffman, Carl A., et al., *Project Dirty Data: Base Level Reporting*, Oklahoma City Air Logistics Center, Tinker AFB, Oklahoma, 15 August 1988.

Gaver, Donald P., Karen E. Isaacson, and John B. Abell, *Estimating Aircraft Recoverable Spares Requirements with Cannibalization of Designated Items*, RAND, R-4213-AF, 1993.

Greenlaw, Colonel Peter G., et al., *Project Dirty Data: Base Level Reporting Follow-Up and Common Item Utilization Study*, Air Force Logistics Command, Wright-Patterson AFB, Ohio, 31 January 1989.

Hodges, James S., with Raymond A. Pyles, *Onward Through the Fog: Uncertainty and Management Adaptations in Systems Analysis and Design*, RAND, R-3760-AF/A/OSD, July 1990.

Isaacson, Karen E., and Patricia M. Boren, *Dyna-METRIC Version 6: An Advanced Capability Assessment Model*, RAND, R-4214-AF, 1993.

King, Randall M., et al., *Improving the Air Force's Computation of Spares Requirements: The Effects of Engines*, Report AF501R5, Logistics Management Institute, Bethesda, MD, December 1986.

Lippiatt, Thomas F., *Variability in the Budget Forecasts for Depot-Level Component Repair*, RAND, N-2930-P&L, 1991.

Miller, Louis W., and John B. Abell, *DRIVE (Distribution and Repair in Variable Environments): Design and Operation of the Ogden Prototype*, RAND, R-4158-AF, 1992.

O'Malley, T. J., *The Aircraft Availability Model: Conceptual Framework and Mathematics*, Logistics Management Institute, Washington, D.C., Task AF 201, June 1983.

Sakulich, Captain Tim, USAF, and Dyann Beatty, *Impacts of Dirty Data (Base Repair and Supply Systems) on Worldwide Requirements*, Air Force Logistics Command, Wright-Patterson AFB, Ohio, February 1989.

Sherbrooke, C. C., *METRIC: A Multi-Echelon Technique for Recoverable Item Control*, RAND, RM-5078-PR, 1966.